

RTX 4080 TESTED MELTING WALLET & SPEED RECORDS

apoc

JANUARY 2023 #515

YOUR EXPERT GUIDE TO TODAY'S TECH




HOW IT ALL BEGAN
THE TRUE STORY
OF THE IBM PC

**STOP YOUR
PHOTOS &
VIDEOS
DECAYING**



**PROTECT YOUR FILES
AGAINST DAMAGE &
LOSS IN QUALITY**



NEW SERIES
**HISTORY OF
PORTABLE
AUDIO PT 1**



CHIP TECH

THE CPU WARS
WHY TSMC
IS SO IMPORTANT

IT'S TIME TO DITCH MICROSOFT!

SECRETS OF THE GOOGLE OFFICE APPS

DISCOVER DOCS & SHEETS TIPS YOU'LL LOVE!

**TALK, DON'T TYPE • MASTER IMAGE USE • TIME-SAVING TIPS
IMPORT & SHARE TO OTHER FORMATS • SAFEGUARD YOUR WORK**

L AU\$10.95 INC GST NZ\$10.90 INC GST
PRINT POST APPROVED: 100000179
F ISSN: 0725-4415



0 1

313006 001362



**INTEL NUC
13 EXTREME**

A TINY LITTLE PC
THAT DOES IT ALL



**SAMSUNG
990 PRO**

QUICKEST SSD WE'VE
EVER TESTED



**LG OLED
ULTRAGEAR
4K SCREEN**

YOU WON'T BELIEVE
YOUR EYES



NBN SMB Solutions for VDSL2 35b & Ethernet WAN (FTTN/B, FTTP, HFC, Fixed Wireless, Satellite)

Vigor2865 Series

Multi-WAN VDSL2 35b Supervectoring & Ethernet routers with 1 x configurable GbE WAN/LAN port, 5 x GbE LAN ports, SPI Firewall, 802.11ax/ac Wi-Fi, VoIP and 32 x VPN tunnels including 16 x SSL-VPN tunnels



Vigor2865ax (802.11ax, WiFi 6)

- Multi-WAN with 1 x VDSL2 35b Supervectoring/ ADSL2+ WAN port, 1 x configurable GbE WAN/LAN port, 2 x USB ports and 2 x Wi-Fi WANs (ac/Vac models)
- 1 x XDSL port and 1 x configurable GbE WAN/LAN port for Failover, Load Balancing and High Availability mode
- 2 x USB 2.0 ports for 3G/4G LTE USB modems, FTP server, network printer or thermometer
- 5 x Gigabit LAN ports with multiple subnets and 60,000 NAT sessions
- 32 x VPN tunnels (including 16 x OpenVPN/ SSL-VPN tunnels) with most security protocols
- Fast VPN throughput, VPN load-balancing and backup for site-to-site applications
- 16 x VLANs for secure and efficient workgroup management
- IPv6 & IPv4; Up to 1022 IP addresses and 8 IP subnets
- Integrated IEEE 802.11ax (WiFi 6)/ac wireless Access Point; dual band; up to 3000 (ax) or 867 (ac/Vac) Mbps throughput
- Object-based SPI Firewall and CSM (Content Security Management) for network security
- VoIP (2 x FXS ports) for Vigor2865Vac
- Support VigorACS 2 and VigorACS 3 Central Management Systems for remote management
- SD WAN capability when used with VigorACS 3
- Central VPN Management for up to 8 remote Vigor routers
- Central AP Management for up to 20 Vigor Access Points
- Central Switch Management for up to 10 VigorSwitches
- 2 years back to base warranty

Vigor2765 Series

VDSL2 35b Supervectoring & Ethernet routers with 1 x configurable GbE WAN/LAN port, 3 x GbE LAN ports, SPI Firewall, 802.11ax/ac Wi-Fi, and 2 x VPN tunnels including 2 x SSL-VPN



Vigor2765ax (802.11ax, Wi Fi 6)

- 1 x VDSL2 35b/ADSL2+ WAN port
- 35b Supervectoring VDSL2
- 1 x configurable GbE WAN/LAN port (P4)
- 3 x Gigabit Ethernet LAN ports with 50,000 NAT sessions
- 2 x USB ports for 3G/4G LTE mobile broadband access, USB printer and storage
- Object-based SPI Firewall, CSM (Content Security Management) and QoS for network security
- IPv6 & IPv4
- 2 x VPN tunnels including 2 x SSL-VPN tunnels
- IPsec VPN throughput at 200 Mbps (AES 256 bits)
- Integrated IEEE 802.11ax (AX3000) wireless Access Point (ax model)
- 2 x FXS VoIP (Vac model)
- Support VigorACS 2 and VigorACS 3 Central Management Systems for remote management
- Central AP Management for up to 2 Vigor Access Points
- 2 years back to base warranty

Vigor167

VDSL2 35b Supervectoring Modem/Router with 2 x GbE LAN ports



- 1 x VDSL2 35b Supervectoring / ADSL2+ WAN port
- 2 x Gigabit Ethernet LAN ports
- Two Operation Modes
 - > VDSL2/ADSL2+ Bridging Mode
 - > Router Mode
- VDSL2 35b download speed up to 300 Mbps
- IPv4 Static Routing
- Supports VigorACS 3 Central Management System for remote management
- 1 year back to base warranty



SUBSCRIPTION ENQUIRIES

(02) 8227 6486
www.techmags.com.au
future@crmaus.com.au

Editorial

Editor: Ben Mansill ben.mansill@futurenet.com
Senior Journalist: Shaun Prescott
Journalist: Joel Burgess
Journalist: Chris Szewczyk
Creative Director: Troy Coleman

Contributors

David Crookes, Tim Danton, Darien Graham-Smith, Kenny Hemphill, Rober Irvine, Nicole Koble, Jeremy Laird, Sam Lewis, Ian Osborne, Les Pounder, Nik Rawlinson, Michael Reed, Becky Scarrott, Chris Szewczyk, Jarred Walton, Mark Williams, Darren Yates

Photography

All copyrights and trademarks are recognised and respected

Advertising

auadvertising@futurenet.com

Management

Managing Director: Neville Daniels
Sales Director: Paul Marttila
paul.marttila@futurenet.com

Printed in Australia by IVE

Distributed in Australia and NZ by Are Direct

ISSN 0725-4415

Future Publishing Australia
PO Box Q1179, Queen Victoria Building, NSW 1230
Corporate Web: www.futurepic.com
Email: apcmag@futurenet.com
Web: www.apcmag.com

We are committed to only using magazine paper which is derived from responsibly managed, certified forestry and chlorine-free manufacture. The paper in this magazine was sourced and produced from sustainable managed forests, conforming to strict environmental and socioeconomic standards. The manufacturing paper mill holds full FSC (Forest Stewardship Council) or PEFC certification and accreditation

All contents © 2018 Future Publishing Australia or published under licence. All rights reserved. No part of this magazine may be used, stored, transmitted or reproduced in any way without the prior written permission of the publisher. Future Publishing Limited (company number 2008885) is registered in England and Wales. Registered office: Quay House, The Ambury, Bath BA1 1UA. All information contained in this publication is for information only and is, as far as we are aware, correct at the time of going to press. Future cannot accept any responsibility for errors or inaccuracies in such information. You are advised to contact manufacturers and retailers directly with regard to the price of products/services referred to in this publication. Apps and websites mentioned in this publication are not under our control. We are not responsible for their contents or any other changes or updates to them. This magazine is fully independent and not affiliated in any way with the companies mentioned herein.

If you submit material to us, you warrant that you own the material and/or have the necessary rights/permissions to supply the material and you automatically grant Future and its licensees a licence to publish your submission in whole or in part in any/all issues and/or editions of publications, in any format published worldwide and on associated websites, social media channels and associated products. Any material you submit is sent at your own risk and, although every care is taken, neither Future nor its employees, agents, subcontractors or licensees shall be liable for loss or damage. We assume all unsolicited material is for publication unless otherwise stated, and reserve the right to edit, amend, adapt all submissions.

Privacy statement

If you provide information about yourself this will be used to provide you with products or services you have requested. We may supply your information to contractors to enable us to do this. Future Publishing Australia will also use your information to inform you of other publications, products, services and events. Future Publishing Australia may also give your information to organisations that are providing special prizes or offers and are clearly associated with the Reader Offer. Unless you tell us not to, Future Publishing Australia may give your information to other organisations that may use it to inform you of other products, services or events. If you would like to gain access to the information Future Publishing Australia holds about you, please contact us.




Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR)
www.futurepic.com

Chief executive Zillah Byng-Thorne
Non-executive chairman Richard Huntingford
Chief financial and strategy officer Penny Larkin-Brand
Tel +44 (0)225 442 244

INTRO

EDITORIAL

WHAT'S INSIDE APC



EDITORIAL

Exiting 2022

With great hopes for 2023



BEN MANSILL
can no longer tell
when one year
ends and another
begins.

Well, I hope you made it through December with few party after effects, your budgets unbruised and at least a few days of down time, with perhaps more to come. My December is just commencing as I write this, but I'm somewhat confident all things will be well.

At the very tail end of a long year it's easy to sit back and take stock, except the gods of tech don't seem to have received the memo about slowing things down...

In this issue we've got some nice treats. Nvidia's RTX 4080 arrived in the Labs, and Chris has once again been running a marathon to test them and share his impressions with you. The performance is, as expected, exactly in the right place relative to the 3000 series and the 4090, but oh my the price, son, the price. Here we are being subjugated by the Australia Tax again. And this time it's a bad one. In the US you can pick up a 4080 for US\$1,199. But at launch the cheapest we could find here was a brutal

"In the US you can pick up a 4080 for US\$1,199. But at launch the cheapest we could find here was a brutal \$2,439. That's outrageous, and a lot of the blame falls on Nvidia and its rather ambitious pricing attitude."

\$2,439. That's just outrageous, and a lot of the blame falls on Nvidia and its rather ambitious pricing attitude. There's no doubt this company is taking the market for every penny it can, despite ongoing whining that it has no choice.

Well, let's see what AMD says about that, as the new Radeons are, we hope, in the back of a courier truck on the way here for testing, and based on the announced US pricing of under a grand, the Australian RRP we hope is somewhat saner.

In brighter news, now that Darren wrapped up his absolutely fantastic Tech that changed Australia – and the world series, he's off and running again with a new series, the first part of which appears in this issue. This time he's

tackling the fascinating history of portable audio, and we're starting with early radio, then moving through all the common ways in which we have enjoyed music and voice over the years, wrapping up with the iPod era. It's off to a great start so do please enjoy this with me, I can't wait to read the rest as Darren pulls it all together masterfully.

And lastly, there's some nice holiday history reading for you in this issue, with the story of the first 'actual' PC, the IBM 5150, as well as the TSMC tale, which is central to our hobby and so interesting. Stay safe you all, and enjoy the summer!

OUR GIFT TO YOU

Exclusive downloads

Your free full-version apps as a thanks for reading APC. These exclusive downloads will only be available for a limited time, from 29/12/22 to 15/02/23.



1

ABELSSOFT CONVERTER4VIDEO


An easy-to-use video converter for Windows.

Converting videos is simpler than ever. Abelssoft proves with Converter4Video that anybody is capable of turning videos from one format into another.

Converter4Video can create files of the proper format for almost any device, directly upload videos to YouTube, and even extract the file's audio track. The video file to be converted is simply dragged and dropped into the Converter4Video window; then, the user designates his desired output format, FLV, WMV, SWF, MPEG, MP4, 3GP, M4V, AVI, MKV, MOV, F4V or RM, and with a click the conversion commences.

Download link:

www.apcmag.com/exclusives



2



ASCOMP IMAGEFORMER

Batch-edit and convert images with ease.

With Image Former you can get your pictures in top shape in no time at all: select any number of pictures, define actions and start conversion. Whether you want to rotate images, crop edges or correct colours, Image Former automatically applies the changes you want to all selected images. The Windows tool also has full control of enlarging or reducing the size and brings your image files into the desired format quickly and easily.

Download link:

www.apcmag.com/exclusives



3


ASHAMPOO WINOPTIMIZER 19

Optimisation software to boost and clean Windows PCs.

Ashampoo WinOptimizer 19 makes Windows PCs faster, cleaner and more discreet. The extensive system analysis provides numerous options to optimise and clean your operating system right from the start page. The built-in cleaner modules have now become genuine classics and rid your machine of even more web browsing traces, temporary files and leftover files from software uninstallations. Auto-Clean, Live-Tuner and Game-Booster ensure your system stays clean, programs launch faster and games run smoothly.

Download link:

www.apcmag.com/exclusives



DISCLAIMER Future Publishing is not the licensor of the exclusive software or any documentation included with it. Rather, Future is merely a distributor of the software, and your use of the software is subject to any accompanying third-party licence terms. You must carefully read and comply with any such third-party licence terms, together with all instructions and README files that come with the software. Accordingly, to the maximum extent permitted by law, all software is provided by Future 'as is' and without warranty either express or implied, and Future will not be liable for any damage that you may incur as a result of using any software downloads. You should back up any important system and data files before using any downloaded software. We recommend that you do not use the exclusive software on a production machine. Nothing in this disclaimer excludes any warranty that may be implied by statute, which may include the Trade Practices Act. While we have taken all reasonable steps to check the software downloads for viruses, we cannot guarantee that it is free from viruses or other harmful code and you should check each download using a virus scanner complete with the latest antivirus updates before use.




INTRODUCING **ALL NEW XL2566K**

360Hz | DyAc+ | XL Setting to Share™ | S-Switch | Shielding Hood

DISCOVER MORE
at **Zowie.BenQ.com**






PAGE 080

IT'S TIME TO DITCH MICROSOFT!

SECRETS OF THE GOOGLE OFFICE APPS

DISCOVER DOCS & SHEETS TIPS YOU'LL LOVE!


TALK, DON'T TYPE • MASTER IMAGE USE • TIME-SAVING TIPS
IMPORT & SHARE TO OTHER FORMATS • SAFEGUARD YOUR WORK



THE TRUE STORY BEHIND THE IBM PERSONAL COMPUTER

The industry-creating IBM Personal
Computer 5150 turned 40 in October

070



HOW TSMC CAME TO RULE THE WORLD


A decade ago, TSMC was generations behind Intel
in chip production technology but now it's the
undisputed world leader

056



GOOGLE PIXEL 7

Properly premium at an appealing price
046



NZXT H7 FLOW

042



STEELSERIES ARCTIS NOVA 7 WIRELESS

044



BENQ ZOWIE 360HZ XL2566K

041



CORSAIR KATAR ELITE WIRELESS

045



SAMSUNG 990 PRO 2TB SSD

A proper contender and 980 Pro successor

048

TECHNOTES

- 010 **News**
What's been going on
- 012 **Tech brief**
Now Intel wants to play nice – with everyone
- 014 **End user**
The gig economy comes under scrutiny
- 015 **Random access**
Intel Evo issue has users pining for MacOS
- 016 **Two bits**
Benchmarking trials and tribulations
- 017 **Tech Talk**
Nvidia's adaptor woes
- 018 **One more thing**
Jon Honeyball calls time on the endless flow of bad news
- 018 **Trade chat**
Launch predictions failed
- 022 **Futures**
Accurate referee calls with high-tech motion capture
- 026 **A-list**
The best products on the market, as picked by our editors
- 029 **The list**
Best gaming PSUs
- 030 **Head to head**
Google Drive vs We Transfer
- 032 **Gadgets**
Techy toys and trinkets

THE LAB

- 034 Razer Blade 15
- 035 Dell Inspiron 16 Plus 7620
- 036 HP Envy x360 13.3-inch 2-in-1
- 037 Intel NUC 13 Extreme
- 038 MSI GeForce RTX 4080 16GB Suprim X
- 039 Asus TUF Gaming RTX 4080 OC Edition
- 040 LG UltraGear 48GQ900 OLED UHD 140Hz
- 041 Benq Zowie 360Hz XL2566K
- 042 Nzxt H7 Flow
- 043 Focal Bathys
- 044 Steel Series Arctic Nova 7
- 048 Samsung 990 Pro 2TB SSD
- 045 Corsair Katar Elite Wireless
- 046 Google Pixel 7a
- 050 Microsoft Modern Mobile Mouse
- 050 Google Pixel Watch
- 051 Draytek Vigor2865ax

SOFTWARE

- 052 VirtualBox 7.0
- 053 Coconut Battery 3 Plus
- 054 Ubuntu 22.101

FEATURES

- 056 How TSMC came to rule the world
- 070 The true story behind the IBM PC
- 076 Understanding IP
- 080 Secrets of the Google Office Apps
- 060 History of portable audio part 1

PC BUILDER

- 085 **System news**
Mark Williams looks on as Nvidia continues to wrangle its product line-up
- 086 **Market watch**
A sampling of PC systems available
- 087 **Blueprints**
Value- and performance-driven hypothetical builds

HOW TO

- 090 **Quick tips**
Solving a variety of tech problems
- 092 **Stop your photos**
and videos decaying
- 094 **Master the Clock**
app
- 096 **Floppy drive-sized**
assembly-based OS

MASTERCLASSES

- 100 **Raspberry Pi**
Arduino to Pi GPIO comms using Python

DOWNTIME

- 102 **Gaming reviews**
High-performance playtime
- 106 **Game changer**
Costume Quest
- 110 **Retro**
Psion Series 5
- 114 **Chip chat**
The less than serious news page



SUBSCRIBE NOW!

Turn to page 20 now
to find out how

HOW WE DO IT

Inside APC

Find out all about APC's editorial policies, test practices, how to read the benchmark results, and more.



"Championing technology doesn't mean we're unrelenting yes-men and -women, however, and APC aims to be as objective as possible in all our coverage."

APC is Australia's oldest consumer technology magazine – having been consistently in print for forty years, since our first issue way back in May 1980 – and we take that heritage and responsibility very seriously. While our focus is obviously on the personal computer the very definition of the PC has changed and shifted markedly since the early 1980s. As such, we touch on many other areas of tech, too, from smartphones and apps to peripherals, accessories, and beyond. We have two goals: to find the best modern tech and to help you make the most of it.

We're also an open church in terms of platforms. We know most people aren't wed to a single brand's products and use a variety of devices. And, like you, APC's journalists want to know what's good in tech – no matter what platform it resides on.

Independent reviews

Championing technology doesn't mean we're unrelenting yes-men and -women, however, and APC aims to be as objective as possible in all our coverage. That means identifying the best products from multiple perspectives – the best performance, best value and best features and, ideally, the products that offer the best mix of these three attributes.

As a matter of policy, reviews published in APC are not shared with product-makers prior to print. We will contact vendors under certain conditions; for example, if we have a problem testing a product that seems to indicate it may be faulty, or to invite a vendor to clarify how a particular feature works. If an APC reviewer has any potential conflicts of interest involving a brand, the review will always be assigned to another writer.

Labs testing

APC strives to conduct the most rigorous, objective scientific tests and benchmarks we can so as to make our reviews as unbiased as possible. We use a variety of tools and programs for this, including many freely available benchmark suites for assessing media encoding, general system performance including storage read and write speeds, gaming and battery life.

In most cases, for the benchmark results published in APC, you can assume that higher is better. There are certain tests that deviate from this rule where the opposite is true; in those cases, we've flagged the results with a note explaining as such.

We use both tables and graphs for displaying results; the latter offers better ease-of-readability, but tables are more compact, so we use these in most cases where thoroughness is preferred.

apc TESTBED

UPDATED
FOR 2023

The current APC testbed used in the Labs for benchmarking all components. This testbed is updated as new and relevant technology comes on stream.

CPU	Intel Core i9 12900K
Motherboard	Asus Maximus Z690 Apex
Graphics Card	Nvidia GeForce RTX 4090 Founders Edition
RAM	2x16GB G.Skill Trident Z5 DDR5-6000
Power Supply	Corsair AX1000
Storage	Seagate Firecuda 530 2TB
Case	Thermaltake Core P8
Cooling	Cooler Master Flux PL360 AIO
Operating system	Windows 11 Pro 22H2

apc AWARDS



APC EDITOR'S CHOICE

When a product scores 4.5 out of 5, it carries the Editor's Choice Award. These are products that exceed expectations and deliver a quality experience up there with the very best.



APC HIGHLY RECOMMENDED

You will see this award if a reviewed product has scored four out of five stars. It means most people can expect satisfying performance from the product, and that we would use it ourselves.



GEFORCE
RTX

MSI GEFORCE RTX 40 SERIES GO BEYOND FAST



CHISELED TO PERFECTION —BUILT TO PERFORM

MSI is proud to introduce a complete line up of custom cards that unleash the full potential of NVIDIA Ada Lovelace architecture.

Featuring extreme silent cooling and the most advanced thermal designs we've ever engineered—MSI GeForce RTX 40 Series graphics cards take gaming and content creating to a whole new level.



SHOP IN-STORE

LEARN MORE
AT [MSI.COM](https://www.msi.com)





SMART HOMES TO GET SMARTER

A common standard for all.

One problem with smart homes is they aren't particularly smart because the IoT devices can't communicate with each other. However, the Connectivity Standards Alliance (formerly Zigbee) has defined a new standard to address this, called Matter. The project started in 2019, and while there have been some delays, the standard is finally finished and over 280 companies have signed up. The result is that certified Matter smart devices will be able to work co-operatively, so you're not tied to a single brand, and don't have to continually update to new standards. It will take some time to work through all the products – right now, there are around 200 certified devices, most of which are light bulbs.



PC GAMING GETS A BIG BOOST

DirectStorage 1.1 allows games to shift decompression from the CPU to the GPU.

DirectStorage 1.1 with GPU decompression is now available on both Windows 11 and Windows 10. The upgraded DirectStorage allows developers to shift the workload of decompressing files from a system's CPU to a GPU. This setup is optimal because the best graphics cards handle parallel processing well. As a result, PCs can load games significantly quicker when DirectStorage 1.1 is implemented. When used in conjunction with a high-speed NVMe drive, game times should be much quicker with DirectStorage 1.1. GPU decompression is available on both Windows 11 and Windows 10, but the former has additional optimisations that make it the recommended OS for the feature. Manufacturers can further optimise hardware to work with DirectStorage 1.1 as well.



LAW

Australia eyeballs new gaming law

Bill would mandate an R18+ rating for any game with gambling-like features.

Australia may be the next country to take legislative action against loot boxes in videogames, as MP Andrew Wilkie has introduced a bill calling for a mandatory adult classification for any game that contains a "gambling-type feature."

Introduced on November 28, the bill says loot boxes come in many forms but share some common characteristics, including that they can offer competitive advantages in games and may also have resale value, which collectively confers on them "significant value for players."

"By tempting players with the potential to win game-changing items, encouraging risk-taking for possible reward, delivering random prizes on an intermittent basis, and encouraging players to keep spending money, loot boxes give rise to many of the same emotions and experiences associated with poker machines and traditional gambling activities," Wilkie's explanatory memorandum says. "This is especially concerning as many games which contain these features are popular with adolescents and young adults."

The debate over whether loot boxes are a form of gambling has been going on for years. Some experts say they absolutely


are, while game publishers and industry bodies have argued strenuously that – from a legal perspective, at least – they are not. The split opinion is perhaps best illustrated in the UK: The mental health director of the National Health Service said in 2020 that loot boxes effectively teach children to gamble, and that same year the House of Lords called for loot boxes to be classified as a form of gambling. But the country's Gambling Commission, which actually makes the

call on these things, has repeatedly ruled that loot boxes are not a form of gambling.

The reason for that position, which has also previously been taken by other bodies including game publishers, rating boards, the Entertainment Software Association, and other gambling regulators, is that items

"Loot boxes give rise to many of the same emotions and experiences associated with poker machines and traditional gambling activities"

drawn from loot boxes have no monetary value. In order to qualify as gambling, the rewards on offer must either be money, or have real-world value supported through official channels. Not every country buys that argument, of course: Both Belgium and the Netherlands have declared that at least some forms of loot boxes violate their gambling laws, and have thus banned them.



Intel's FakeCatcher

What even is reality, anymore?

Deepfakes, videos doctored by AI to appear real when they're not, have been around for years. But they hit the headlines around 2018 when it became possible to create them on modest systems with readily available software. There was something of a panic as it became clear that our world of digital imagery was forever untrustworthy. Editing images is as old as the invention of photography, but deepfakes made realistic videos a reality.

At their simplest, deepfakes work by mapping an AI-manipulated face over the movements and expressions of an actor to alter an original clip. Intel's

FakeCatcher software has a claimed accuracy of 96 percent and can spot a fake in milliseconds. It can even uncover the source model with a 93 percent accuracy. How does it manage this? It uses blood. As your blood circulates around your body, your veins pulse and change colour, an effect called photoplethysmography (PPG). FakeCatcher looks for PPG signals across the face and creates a map. This is run through some machine learning wizardry to see if it matches reality. Intel also uses eye gaze detection. Fakes and fake detectors look set to battle it out for the foreseeable future.

ZEN 4'S V-CACHE CHIP DUE SOON

But they will be limited to eight-core.

AMD's secret weapon for the Zen 3 range was its V-cache version, the Ryzen 7 5800X3D. It stacked high-speed memory on top of eight cores and tripled the L3 cache, making for the world's fastest gaming chip on launch. Although Intel responded with the Core i9-12900KS to steal some of that thunder, Intel's chip was faster in some games, albeit more expensive. AMD had no trouble selling its V-cache chips, and even now they're still hard to find. The same V-cache editions will be released for the Zen 4 Ryzens and should be due in the first half of 2023. The hope was that the cache trick would be applied to a higher-end version with more cores, but rumour has it that AMD will start small, with a six-core Ryzen 5 7600X3D and an eight-core Ryzen 7 7700X3D.



FASTEST HARD DRIVES EVER

Matches SATA SSD speeds.

Seagate's second generation of Mach.2 drives, the Exos 2X18 and 2X16, can manage data transfer rates of 545MB/s over a SATA 3 interface, which is comparable to SATA SSD performance. They build on the slower 2X14 that began shipping in 2019 and, while they're no match for an M.2 NVMe SSD, these mechanical drives come in larger capacities of 16TB and 18TB. The magic isn't achieved by spin rates – those are a modest 7,200rpm – but using two actuators effectively means it can act as two drives simultaneously. One area where they won't match solid state is in seek times, which at 4.16ms are nowhere near the microseconds of an SSD. These drives are aimed at server work and 'hyperscale' workloads, so are robustly built, with SAS and self-encrypting versions too.

NEW DATA TRANSFER RECORD

The internet on one fibre.

A new data transfer record using a single strand of a fibre optic cable has been set at 1.53 petabits. A petabit is 1,000 terabits. To put this into perspective, the latest study by TeleGeography, an international telecommunications consulting company, estimates that total internet traffic for 2022 averages at 997Tbps, under one petabit, so you could run the entire global internet down one 0.125mm fibre.

The feat was achieved by a team at the National Institute of Information and Communications Technology in Japan. They used light in 55 specific frequencies, a process known as multiflexing. It's not the fastest data transfer (1.84 petabits on an experimental optical chip), but it used standard fibre, so the existing infrastructure could be upgraded. ■

TECH BRIEF

Now Intel wants to play nice with everyone

Chipzilla shows a sudden willingness to share with rivals.

For decades Intel's relationship with chipmaking rivals such as AMD and Nvidia has been a feisty one, often seeking to exploit its domination at all costs.

In 1999 for example, the company got into a tangle with America's Federal Trade Commission over how it was protecting technical information and patents. And a decade later, the European

Commission found "abuse of dominance" as Intel paid retailers and computer manufacturers to only use Intel CPUs.

So it's perhaps a little unexpected that Intel is now

suddenly singing kumbaya and effectively saying to its rivals "can't we all just get along?"

"We should have an open programming language... that can run on Intel, AMD and Nvidia GPUs without changing your code," Intel CTO Greg Lavender told the business website Venturebeat, as he sought to promote the company's investments in SYCL

(pronounced 'sickle'), a Java-like programming language that bridges the gap between different CPU and GPU hardware.

It's a language designed for software where processes are "heterogeneous" – tasks such as processing large amounts of data in parallel where you might throw a mass of GPUs at the problem. As such, it's already used by various scientific supercomputers and in software such as GROMACS, which scientists use to simulate the behaviour of proteins, and which came in very handy during the early days of the pandemic.

What makes SYCL a significant break for Intel is that it is entirely open, and is by definition designed to make the same code work across different GPUs and CPUs. In theory, you could write the code and have it run on everything from an Intel i3 to an AMD Ryzen or an Nvidia GeForce GTX.

And to really show it meant business, earlier this year Intel bought Codeplay, a company that has contributed extensively to SYCL, as part of what it calls its "open standards vision".

"Writing software is really expensive," said Dr Tom Deakin,

a lecturer in advanced computer systems, and chair of the SYCL advisory panel, who argues that given the choice, people want to write portable code that will run on different hardware.

"If they're going to have that code live for a long time, they can't afford to keep rewriting it."

SYCL is not the first programming language to attempt to bridge the gap between different types of hardware, with the likes of OpenCL (which is managed by Kronos, the same group who run SYCL), but what makes it different is the level of abstraction – essentially taking care of high-level processes, making it quicker and easier to write code.


Deakin argues that this is good for coders, as it means there is no risk of having to rely on the fortunes and decisions of only one company. "There's lots of compiler support for it and lots of those compilers support multiple devices from multiple vendors," said Deakin. "I always like to think this multiple compiler position is the sign of a healthy ecosystem."

And this perhaps reveals why pushing SYCL isn't an entirely altruistic move from Intel. The company has a close to 70 percent share of the GPU market if you include integrated graphics chips on devices such as laptops, but it controls only a tiny proportion of the discrete GPU market, which is dominated by Nvidia. SYCL could be an attempt to challenge Nvidia's domination of GPUs.

The move also aligns with Intel's broader strategic shifts. Intel is in the process of reorganising its business away from controlling the entire chip ecosystem, to a more mercenary approach, where it will build whatever chips customers want.

Nicole Kobie ■

"We should have an open programming language... that can run on Intel, AMD and Nvidia GPUs without changing your code"



BILLION

5GHz 1733Mbps
2.4GHz 600Mbps



CAT.6

BiPAC 4520VAOZ R3



First International
Trading Pty Ltd.

4G/LTE Dual-SIM Dual-Band Wireless VoIP VPN Router



Enables high-speed and secured Internet connectivity for home, SOHO, and enterprise users.

- ▶ 4G LTE Mobility
- ▶ Dual Radio/Dual SIMs
- ▶ 4X4(1733Mbps at 5GHz)/3x3(600Mbps at 2.4GHz)
- ▶ Automatic Failover/Failback between 4G/LTE and Gigabit EWAN



au.billion.com

(02) 8318 3777

sales@firstint.com.au

Unit 5, 1 Talavera Road, Macquarie Park NSW 2113

END USER

Food delivery flounders as the gig economy comes under scrutiny

The days of Uber-driven exploitation may be numbered in Australia, writes Shaun Prescott.



Shaun is an Australian editor for *PC Gamer*, *GamesRadar* and *PLAY*, and writes for *APC*, *TechRadar* and more.



When Deliveroo announced it was exiting the Australian market in November, it was difficult for critics to resist a smirk. Part of the ever-expanding gig economy model, where self-employed independent contractors work on a piecemeal basis, it seemed to presage a larger contraction in that notoriously precarious employment model. Like Uber and Amazon, Deliveroo didn't employ food delivery drivers per se, but instead used a smartphone app which workers, on a self-employed contractor basis, would log into to pick up jobs and shifts.

At first glance it's an appealingly flexible system, but the reality is bleak. In Deliveroo's case, workers were beholden to an algorithmic approach to their work suitability. If, for example, they cancelled a shift too many times, their selection for future shifts was algorithmically rendered less likely. Those who chose to work during peak times were reportedly preferred, too. In other words, the vaguely utopian idea of a kind of cash injection style of employment – just log in and work when you want – can easily become a very precarious full-time job for those who, for example, suddenly lost their jobs at the height of the COVID-19 pandemic.

Anyway, upon closer inspection there's nothing utopian about it; it's another example of tech corporations side-skirting employment responsibilities – rights, holidays, sick leave etc – in order to save their own bottom line. When the market shrinks, to name one example of this model's benefits to businesses, there are no redundancies to pay out. It also means there is no minimum wage, no superannuation, and no unfair dismissal protection.

Deliveroo's exit from the market doesn't signal Uber's exit from the market – the biggest and gravest offender in the gig economy world. In the weeks following Deliveroo's exit, Menulog – which has always enjoyed a much bigger chunk of the market over Deliveroo – cut “a small number of roles,” a spokesperson told the SMH. Whether that is connected with a wider post-pandemic slump in the market, or to the company's move last year to employ some delivery drivers on a full-time basis, it is hard to know – the company itself is not forthcoming. DoorDash also cut 1,250 staff in early December, including workers in its Australian office.

More likely, is that the

Albanese government's plans to scrutinise the gig economy model are looming. Next year legislation is set to be introduced which will focus on improving conditions for gig workers. In late August, workplace relations minister Tony Burke said in a speech (via the SMH) that “gig work drives down wages and it has been spreading like a cancer through the economy,” which is certainly not a favourable view of the model. He went on: “Gig work must not become the equivalent of a get-out-of-jail free card in Monopoly where business is able to avoid the minimum standards that Australians hold dear – 21st century technology must not mean 19th century working conditions.”

And action is coming: after sweeping industrial relation reforms in early December, a second round is forthcoming in 2023, and the gig economy is firmly in the government's targets. Little wonder then, that Deliveroo is dead, Menulog and DoorDash is shrinking, and Uber, presumably, is watching with interest. The best case scenario? We can still order pad kee mao from our phone in front of M*A*S*H reruns, but the person who delivers it is fairly employed.

Intel Evo issue has users pining for MacOS

What's better than a laptop that can't hold its charge?

As someone who has to set up at least a few laptops every week, I have a long list of personal grievances with Windows 11, but there's one 'feature' that is way more of a deal breaker than anything else we've seen.

Occasionally, modern laptops that are supposed to be asleep seem to get really hot and heavily drain the battery. We'd noticed this issue on some of the many laptops that come through APC's doors over the last couple of years, but wrote it off as a specific system glitches or abnormalities arising from some of the system settings we change in order to benchmark performance. The team over at *Linus Tech Tips*, however, did some digging and found there's been a pretty big change in the way Intel computers sleep ever since the launch of instant wake features, and it means that when you put your laptop to sleep fully charged, there's no guarantee it'll have any battery left when you next open it up.

One of the features Intel wanted to introduce with the 11th Gen Evo laptops was a system that would be able to update itself in the background and, according to *Linus Tech Tips*, Intel changed the default sleep state settings from S3 – which only left power running to the memory modules – to a new S0 state that allowed the processor to continue to function if needed. This Modern Standby state also allows the device to remain connected to the network, which means it can update the system, if needed.

When a laptop is plugged in, it makes a lot of sense for this feature to run whatever

updates are needed in the background to save you having to do it when you need to use the computer, but it seems that some systems are performing these background updates when the computer is not connected to a power source as well. For some users this will be fine, but for anyone who has charged their laptop and needs it to maintain that charge while sleeping this isn't an ideal solution.

Unfortunately many laptop vendors have not included support for both S3 and S0 sleep modes in their firmware, which means there's not necessarily a quick fix you can implement through Windows (unlike simple settings options that fix this issue on MacOS). You can check if you are one of a handful of lucky devices to offer S3 sleep by going to your command line and typing `powercfg /a` and seeing if the system supports Standby (S3). If it does, you should have a BIOS setting

that'll allow you to disable network connected sleep... if S3 sleep is not supported then you'll likely have to settle for a different workaround, or none at all.

For the many users that won't have access to S3 sleep settings, you may want to go into your Power and Sleep settings and change the 'When I close the lid' option to Hibernate, as long as you're ok for much slower boot times.

If you don't want to have to sacrifice startup times for battery life, then there's one other possible option: remove the power cable well before you put your laptop to sleep... as this should disconnect network connectivity in S0. Having your computer connected to power, closing the lid and then disconnecting may actually be tricking your laptop into staying in a network-connected sleep mode that can heavily drain the device.



Joel is one of the senior journalists on APC magazine and apart from being the resident laptop, monitor and fitness tech geek, he's also pretty into surfing his local Sydney beach breaks when he's not in front of a screen.



"It makes a lot of sense for this feature to run whatever updates are needed in the background to save you having to do it when you need to use the computer, but it seems that some systems are performing these background updates when the computer is not connected to a power source as well."

TWO BITS

The trials and tribulations of PC benchmarking

Oops, I'd better retest that...



Chris has been elbow deep in PCs since before the turn of the millennium, and he wouldn't have it any other way.

When we judge a PC, laptop or component, we obviously rely on benchmarks. Other than price, it's the primary differentiating factor between any two systems or parts. In order for us to bring you the best analysis, reviews and interpretations, it goes without saying that there's a requirement for accurate data. Most of the time, results fall within expectations, but other times, they do not. Just recently I experienced one of those 'oh no' moments.

I have recently migrated to a new test platform consisting of an Intel Core i9-12900K and Asus Z690 motherboard with Windows 11 22H2. It's a high quality system. Prior to the launch of the Nvidia RTX 4090, I embarked on the very time-consuming quest to retest several graphics cards on this new platform. All was well with the world.

After APC issue 514 went to print, I had some discussions with colleagues. It turns out that many of my results were outside of expectations. Some of my RTX 30 results were over 10 percent lower than where they should have been. I was aware of a potential issue with the GeForce Experience software, but I never install that so I paid it no heed.

Subsequent research led me to the information that Windows 11 22H2 can produce poor game performance. It's cause is attributed to the activation of otherwise hidden debugging tools that degrade performance.

In the end I was partially lucky because the Nvidia cards were consistently affected, with no real outliers. The final conclusion would not have changed. It's not like a 4090 was suddenly slower than a 3080 Ti or anything. In most cases the AMD RX 6000 cards were unaffected. Also, because I have switched over to the use of several heavily ray traced benchmarks, the AMD cards already looked quite poor when compared to their Nvidia competitors. Of course, with the imminent release of the RX 7900 XT and XTX, those ray tracing results will become much more relevant.

I apologise to our readers for this slip up. Thankfully, Microsoft has reportedly resolved the issue in its Nov 29th KB5020044 Windows 11 update. The GPU benchmarks in this issue have been updated with the correct numbers.

Having spoken about relying on data to draw conclusions, allow me to throw that out the

window and speculate a little on just where AMD's RX 7900 XT and 7900 XTX cards will end up. Without having tested the cards myself yet, it looks like the 7900 XTX will be very competitive with the RTX 4080 in traditional raster-based titles. Once ray tracing and DLSS enter the equation, the 4080 will probably take the lead. That leads me to another benchmarking conundrum.

AMD and Nvidia are putting ever increasing emphasis on their upscaling technologies, DLSS and FSR. The increasing prevalence of these tools, while fantastic for the end user, makes apples to apples comparisons difficult. Watch out for AMD and Nvidia slides, read the fine print and also the scale used in graphs. Some benchmark results are not what they seem.

Stay tuned for our next issue. We'll have reviews of the 7900 XT and 7900 XTX. You can be sure that I'll be taking extra care to ensure that results from these cards (and the Nvidia competitors) will be in line with expectations. Accuracy, fairness and impartiality are tough standards to live up to if your benchmark data is unreliable! ■

Chris Szewczyk

Nvidia facing serious 4090 cable issues

Ongoing RTX 4090 16-pin adapter woes.

Last month, we finally got a taste of what Nvidia's Ada Lovelace architecture and the RTX 4090 GPU have to offer. Yet already there are reports of failing cards, which we touched on in the last issue of APC. It's getting serious, because for the second major architecture launch in a row, Nvidia's greatest graphics cards are raising power-related concerns.

That's not particularly surprising, given the new king of GPUs has a base power draw of 450W and custom cards can push that up to nearly 600W. But there's a potentially bigger problem, and it involves Nvidia's 'quadropus' 8-pin to 16-pin adapter cable.

All RTX 4090 cards appear to come with the same adapter cable, bearing an Nvidia logo. On one side is the new 16-pin 12VHPWR ATX 3.0 PCIe 5.0 connector and, on the other, are four traditional 8-pin PEG connectors, each capable of providing up to 150W of power – potentially up to 600W of power through a single cable. How convenient!

Except when you look at the 16-pin connector and compare it with the existing 8-pin connector, you'll notice that the new connector is physically about two-thirds the size of the old connector. Think about that for a moment – four times the power coming through a smaller physical connection. Presumably, PCI-SIG, the ruling body that helps to define PCI-related standards, knows what it's doing. Maybe.

The problem is that the 16-pin connector carries so much current and power that if it doesn't make a secure connection, things could go wrong. Or at least that's one potential theory. So far, there have been at least a few reports of melted connectors posted online. With potentially thousands or even tens of thousands of RTX 4090 cards in the hands of users, there are bound to be a few problems. What's unclear right now is how

widespread the problem is, though Nvidia is aware of the reports and is in contact with the owners.

CableMod, which makes custom-sleeved cables, has some guidance on the 12VHPWR connector. "It appears that bending the wires too close to the connector could result in some of the terminals coming loose or misaligning within the connector itself. This may lead to an uneven load across the other wires, increasing the risk of overheating damage. This is substantially higher if the bend is horizontal in relation to the connector orientation (left to right)." It recommends only bending cables starting at 35mm away from the connector. Considering the size of many RTX 4090 cards, taking that advice may prove impossible in many cases.

Pictures posted online so far

have shown some 'horizontal bend' near the connector. This wouldn't normally be a problem, but it's possible the extra stress on the connection could cause the pins inside the connectors to become loose. If that happens, electrical arcing could occur, which could then cause heat and melted plastic.

This isn't the first time Nvidia has had launch issues with a new halo product. Previous-generation Ampere RTX 3080 and RTX 3090 cards experienced instability at launch – some models had problems so severe that a recall and replacement program had to be instituted.


For now, our advice is to ensure your 16-pin connection is secure and avoid any stress on the cables. If you can find a PSU with a native 16-pin connection or a right-angle adapter, that may help.

Jarrod Walton ■



Jarrod Walton
Jarrod can barely contain his excitement with all these big launches.

CableMod suggests not bending the adapter closer than 35mm from the connector. Not easy here!



TRADE CHAT

The tech drama season

That prediction didn't work out so well...



Jeremy Laird
Six raw 4K panels
for breakfast, laced
with extract of
x86... Jeremy Laird
eats and breathes
PC technology.

"I suspect both Raptor Lake and its Alder Lake predecessor were originally intended for Intel's 7nm node rather than 10nm."

Roughly a year ago, I predicted 18 months of awesomeness for the PC. "A rapid-fire product assault is coming from AMD, Intel, and Nvidia, the likes of which we've rarely seen," were my exact words. With the reveals under our belts, it hasn't worked so well. Am I disappointed? Yup. But let's keep some perspective.

The five launches are, of course, new CPU families from Intel and AMD respectively, graphics cards from AMD and Nvidia, and Intel's long-awaited entry into the gaming GPU market. As I write

this, it's been pretty patchy so far.

So, what happened – apart from me simply being wrong? There hasn't been a single overarching

theme. It's a mix of failure to execute, unrealistic expectations, unusual market conditions, and cynical production planning. It's just unfortunate that it has all come together at the same time.

The failure to execute obviously applies to the dumpster fire that is Intel's Arc graphics launch. Had those GPUs rolled out on time, they would have been merely

disappointing. But arriving a year late has compounded Arc's woes. The hardware isn't awful, which gives me hope for the future. But Intel needs to get a grip of its graphics drivers to have any hope of competing with AMD and Nvidia.

As for the unrealistic expectations, those apply to AMD's Zen 4 architecture and the Ryzen 7000 CPU family. In terms of headline features, including 5nm production tech and 5GHz-plus clocks, Zen 4 delivers. But its underlying architecture is only a minor tweak compared to Zen 3. And AMD didn't up the core count. So, it's an incremental improvement rather than a revolutionary step forward.

Moving on to Nvidia's new GeForce RTX 40 graphics cards, Nvidia pretty much delivered on my expectation that certain SKUs would be capable of double the performance of the RTX 3000 series. The new RTX 4090 is a monster. But Nvidia positively screwed the pooch with the RTX 4080 boards and I await further horrors with the rest of the 40 Series. That's at least in part down to market conditions. Nvidia has a mountain of old 30 series graphics cards it needs to shift, just as demand for GPUs is cratering thanks to a sudden implosion of


crypto mining demand and an increasingly sickly world economy.

Of all the launches so far, Intel's Raptor Lake has probably come closest to meeting expectations. By ramping up both the core counts and clock speeds, Raptor Lake packs a serious punch. But I'm uneasy about the power consumption.

I suspect both Raptor Lake and its Alder Lake predecessor were originally intended for Intel's 7nm node rather than 10nm. Consequently, they're both hot and hungry architectures. That's a problem I'm not expecting to disappear when Meteor Lake arrives on 7nm next year. If Alder Lake and Raptor Lake were meant to be on 7nm, Meteor Lake was probably intended for 4nm. It's going to take a while before Intel's production nodes and architectures fully align again.

If all that makes for a downbeat read, here's the thing. The RTX 4080 is annoying and Intel messed up Arc. However, there's still some great new technology available for the PC and plenty more to come soon. If you love the PC, it's too easy to get carried away by product launches and road maps. If that's my bad, the immediate future of the PC is far from gloomy. It's just not quite as stellar as I'd hoped.


Jeremy Laird ■



ONE MORE THING

Turning off the tap

Jon Honeyball calls time on the endless flow of bad news.



How much is too much? Obviously, I'm not referring to gin. Or a nice wine. Or even the smile of someone you have helped.

No, I am referring to the flood of news that seems to be overtaking everything in our lives. Maybe the start of Covid, and the work-from-home mantra, accelerated the dreaded FOMA or "fear of missing out". Certainly it prompted us all to pay unprecedented levels of attention to the big questions: how was Covid progressing? Which form of lockdown were we in? Could we speak to five people over the garden fence or was it only three?

I think there was a small, subtle, but important shift in the way we consume news from that point. No longer would we wait for news media to deliver their latest pontifications. Now, we could get the self-same commentators via Twitter delivering their missives in usefully short soundbites. Forget any kind of context or analysis – you could be whisked away to a long-form page of text on some website if required, complete with adverts. Meanwhile the spoken format of podcasts seemed to gain relevance, at least in my circles. Maybe it was because you could listen to this

stuff while driving to work, and have a change from the same radio commentators repeating the same thing on a ten-minute rotation cycle.

With the arrival of 2022, we all hoped for a return to normality, whatever that was. But no, this was not to be: the unfolding war in Ukraine led to a growing worry about the economy, energy prices and where to buy petrol at a decent price. Or even at all. Then Her Majesty died. The news cycles reached unknown levels of hysteria.

So surely we have reached Peak News Flood? Of course not. But now I have reached peak exhaustion. I am clamping down on my news feeds, and trying to Take Back Control (remember that?). And yet the news flood keeps increasing. This morning I went down a rabbit hole reading an expert's article explaining how the population of snow crabs has been decimated. All to do with arctic ice shelves, and what happens deep in the ocean.

The more you hear and read, the worse things appear to be. And there is actually nothing that can be done by the likes of you and me. Of course, we can turn the heating down, think about the replacement cycle for the car and decide that it will be electric when it finally dies.

Small changes are good, even if they are small, as they add up. But, despite all the good intentions and worthy efforts, it is so dispiriting when being flooded by all of this noise as, deep in our hearts, we know that our efforts are irrelevant in the greater picture.

Maybe it is time to cut off the tap. To walk away from all the ya-hoo-nonsense of the politicians and political commentators. To continue to make donations to appropriate charities that can do real good out there, pull the lid over my head and then weld it shut from the inside.

Little wonder that many people watch YouTube videos of cats to calm them down, but I am yet to find my equivalent. Perhaps a channel showing wood logs burning in a fireplace, with the subtle crackle of the embers? Except I would then feel guilty about the carbon generation.

Jon Honeyball ■



Jon Honeyball
Jon would love to receive junk email about this subject. Simply send it to billg@microsoft.com.

"So surely we have reached Peak News Flood? Of course not. But now I have reached peak exhaustion. I am clamping down on my news feeds, and trying to Take Back Control (remember that?)."

SPECIAL
OFFER



Simply scan the QR
code above with your
smartphone camera



SPECIAL OFFER

Subscribe to APC

The go-to resource for PC enthusiasts

All the convenience of *APC* delivered
direct to your door. Get the print edition
home delivered for a year for \$79

Save up to 47% off the newsstand price Get the print
edition home delivered 24 issues only \$139.00
– that's just \$5.79 per issue!

SPECIAL OFFER

APC digital edition also available

Subscriptions also available for Apple
and Android devices.

APPLE OFFER

Free 30 day APC
app subscription trial.
Read our latest issue
free today



The Apple edition for iPhone,
iPad and iPod touch
tinyurl.com/APCAPPLESTORE
From \$4.99 an issue

The Zinio edition for all
Android devices
apcmag.com
From \$4.99 an issue

Two easy ways to subscribe to APC magazine:

VISIT www.techmags.com.au or CALL (02) 8227 6486



SCAN ME

TERMS & CONDITIONS: This offer expires January 25th 2023, and is available to subscribers within Australia only. Please allow 6-8 weeks for delivery of your first issue. For full terms and conditions, see: <http://techmags.com.au/style/images/FutureTC.pdf> Please see our privacy policy on page 3. *Six issues for \$49, then \$49 charged to your credit or debit card every six issues thereafter on auto-renewal.playback); 32x22x1.7cm; 1.78kg.

Settling the offside debate... with science!

One researcher sought to settle whether cameras help offside calls using high-tech motion capture. Nicole Kobie reveals the final decision.



Tottenham's Son Heung-Min's cross - with the slightest brush of a teammate's head - swung into the goal ten minutes into a match against Fulham in September, but VAR revealed Harry Kane was offside and the goal was disallowed. Yet more drama unfolded in the 90th minute, when a Spurs goal was again disallowed after a VAR review.

Welcome to the wild world of video assistant referees (VAR). While those calls might have been correct, recent controversies prove that's not always the case - but AI, 3D motion capture and virtual reality could all help.

The Premier League took its time before introducing technology to aid referees with line and goal calls, but finally introduced VAR systems in the 2019-20 season for some offside, goals and penalties. That was welcomed by some, derided by others. Amid a series of controversial video calls that year, former footballer and broadcaster Gary Lineker said VAR was "killing the game". However, the Premier League

claims match-decision accuracy increased from 82 percent to 94 percent thanks to the technology, and research shows the number of offside calls and fouls has decreased since its introduction.

To address the challenge, FIFA and UEFA are both trialling semi-automated offside technology to remove some of the subjectivity from offside call rulings and to speed up decisions. The controversy over VAR calls highlights the importance of introducing technology to sport carefully. "Technology can have an immediate and noticeable effect on sport like VAR has had, where it is clear to see changes in the game," said Jonathan Taylor, a sports scientist at the Teesside University.

But with the controversy over VAR continuing, improvements are clearly needed. In September, referees' organisation the Professional Game Match Officials Board said it would "fully cooperate and collaborate" with an investigation into two goals that were disallowed using VAR at the Chelsea-West Ham



Spurs had two goals disallowed by VAR in their match against Fulham.

and Newcastle-Crystal Palace matches in September, effectively admitting the calls were incorrect.

VAR challenges

So is VAR fit for purpose - and can it be improved? Pooya Soltani decided to find out for himself. A researcher at the University of Bath's Innovation Studio, Soltani is an expert in applying 3D motion capture to sports. Last year, while scrolling Twitter, he spotted a debate over an offside incident in a match between Brighton and Liverpool, and figured he could answer the question. And that gave him unique skills to unpick this problem of whether VAR helps or hinders referees in football.

"Our Innovation Studio is all about cameras," he explained. "We use a variety of tools and techniques to study motion in humans and animals."

VAR can help football officials, but we should be aware of its limitations, Soltani says, explaining that in order to monitor the offside rule, referees need to monitor two

key events: first, when a player passes the ball to a teammate; and second, whether that teammate is offside. In particular, at the point of the pass, is the receiving player in the opposing team's half and closer to the goal line than the ball and the last outfield player?

This is not easy to monitor in athletes running full pelt from different directions. "The players might be well apart, which makes it even harder to keep an eye on the ball-kick moment and the players in [an] offside [position]," Soltani said. "[Referees'] line of vision could also be obscured by other players."


That's where VAR is helpful: it lends an extra pair of eyes to referees. "When in doubt, they can review the event from different angles and at lower playback speeds," he said. "The second aspect is particularly important when we are dealing with fast movement interactions with the ball where our vision is simply not able to see the critical moments."

But there are technical limitations, and referees should be aware that the frame rate of broadcast cameras is usually about 50fps. "The problem happens when they pause the video in critical moments," Soltani explained. "But players' actions are faster than this and when you pause the video, their bodies could appear blurry."

Such complications may partially explain the Premier League's delay in adopting the technology, says Taylor, saying it is a much more complex intervention in the sport than goal-line technology, which was introduced to the league in 2013. "How successful the integration of VAR has been is debatable," he said. "For instance, while it can and has enhanced the accuracy of a lot of the decisions made, there are still issues with offside decisions."

Taking a closer look

To find out how accurate the technology is, Soltani set up a study in the Innovation Studio using motion-capture cameras, digital analytics and a set of football-loving volunteers. Motion capture has been used for biomechanics tracking since the 1970s, says Soltani, for everything from computer



Above: Killing the game? This is a familiar sight at Premier League grounds.

animation and video games to military research. And, of course, movies – think of performers such as Andy Serkis in a green suit to play Gollum in the *Lord of the Rings*.

The technology has moved on somewhat since the days of disco, however. "You place reflective markers on anatomical landmarks of the performer," Soltani explained. "The system uses data from image sensors to triangulate the 3D position of a marker between two or more cameras to provide overlapping projections."

For this project, Soltani set up a series of offside scenarios, asking volunteers to play them out while he recorded their passes from multiple viewing angles. The ball and players all had reflective markers to let Soltani record their positions in 3D. "One of the players passed the ball to his teammate who was standing next to an opponent," he said. "I then showed the video to ten volunteers and asked them

to identify the ball-kick moment and to judge whether the other player was offside. I compared their answer with the data from the motion-capture system."

How accurate were those video calls? "I showed that when using the VAR system, participants identified the ball-kick later than the actual moment," he says, pinning the delay at 132 milliseconds. "This delay could be very small – in milliseconds – but in fast-paced sports like football, this delay could put players in another location. For example, if a player moves at around eight metres per second, a delay of 132 milliseconds could correspond to one metre."

The study also revealed the best places to put the cameras. "In terms of camera angles, I also showed that when the cameras are positioned on the sideline and goal line, participants have more accurate judgements," he said. "From the corner view, and when the attacker was on the left side of the defender, sometimes the attacker appeared to be closer to the goal line, resulting in wrong offside decisions."

Taylor says the study helps explain the practical limitations of VAR and offers solid solutions. "Strategies that aim to enhance the process need to focus primarily on making the footage as clear as possible for VAR [officials] to make their decision," he added. "Unfortunately, there isn't much that can be done around potential human error other

Below: Games outside the Premier League still rely on human referees' assistants.



than training and ongoing review of decisions, which I believe is common practice."

Making better calls

While the study was small, with only 12 participants, it offers ways that football leagues could improve the accuracy of their VAR-assisted calls. First, ensuring cameras are best placed for good angles, and second, boosting their frame rate to better detect kicks when watched in slow motion. "This way, when you pause the video, the images won't appear as blurry as using normal cameras," Taylor said.

Some credit should go to the international governing body, FIFA, he says, as it has begun introducing improvements for marginal decisions. "Instead of using a thin line for marking the players, they now use thicker lines to compensate for uncertainty," he said.

And FIFA does have its own test procedure for VAR, which will look at quality, latency and synchronicity of the video feed, says Soltani. "They have also introduced AI to drive computer avatars based on 29 data points extracted from the players," he explained. "We have to see how the new technology will perform and what issues can arise in terms of communicating the data, synchronisation of ball data with video, calibration of the system and other controversies."

Alongside the automation being added to VAR, there are other cutting-edge technologies that could help, he adds. "Alternative technologies like photogrammetry that extracts 3D measurements from 2D data (images) could be used," Soltani said. "The images could be merged to create a 3D avatar of the subjects, and when viewed on a 3D display or a headset, it can be rotated in any direction." And for the next stage of his research, Soltani is looking at how contextual information from the field can affect VAR operators' perception.

But even with motion capture, 3D avatars and the new automation assistance, it's worth remembering that a human referee still makes the final decision – these tools don't replace officials, but aim to help them do their jobs better. "VAR is definitely an improvement to the game but controversies may still arise because the outcome is still dependent on human interpretation," Soltani said.

So don't worry: even though football officiating technology is coming on in leaps and bounds, you'll still be able to scream at the ref from the stands and argue with your mates in the pub. ■

AUTOMATING OFFSIDE

As we've seen, using VAR to judge offside calls has two complaints: poor viewing angles and frame rates lead to bad calls, and the analysis can take too long. A solution is on the way: semi-automated offside technology (SAOT).

Three years in the making, the SAOT system pairs artificial intelligence, real-time sensors and smart cameras to better visualise the position of the ball and players. A sensor known as an inertial measurement unit will be tucked inside match balls, recording data 500 times a second and sending it back to the operations room, with roof-mounted cameras also keeping watch on its position. Those cameras will track more than two dozen points on each player – in particular, swinging limbs that are relevant to offside calls – updating 50 times a second for perfect pitch position.

If SAOT spots a potential infraction, the AI-powered system will automatically alert officials in the video operation room, who will quickly check the play manually. The hope is that the automation will reduce analysis time from 70 seconds on average to just 25.

As with VAR, the ultimate decision is in the hands of human referees, who will stay in the loop by validating the automated system's conclusion before relaying the ruling to on-field officials. If officials disagree, they can manually select the kick point and use existing systems to show where the offside line should lie. Indeed, FIFA has described SAOT as a "support tool for the video match officials and the on-field officials to help them make faster, more accurate and more reproducible offside decisions".

Fans will then be shown an




Semi-automated offside technology will be used in the World Cup and could make VAR decisions much quicker.

automatically generated 3D animation of the incident, either on the big screens in stadiums or via broadcast partners on TVs at home or in the pub, letting spectators see what officials do.

Pierluigi Collina, chairman of the FIFA Referees Committee, said in a statement over the summer that VAR had reduced the number of offside mistakes. "We expect that semi-automated offside technology can take us a step further," Collina said at the time. "We are aware that sometimes the process to check a possible offside takes too long, especially when the offside incident is very tight. This is where semi-automated offside technology comes in – to offer faster and more accurate decisions."

The technology is already approved by FIFA and is set to be used in the World Cup in November, following trials in lower-level leagues. In February, Zayed Al Ameri of Al Jazeera was the first player to have a goal disallowed following a SAOT review in the Club World Cup in Qatar. UEFA has also trialled the system. ■



Cameras will track more than two dozen points on each player as well as a sensor inside the ball.



The source for tech buying advice

techradar.com

THE A-LIST

The A-List

The best products on the market, as picked by our editors.

FIND THE BEST DEALS ON NEW GEAR AT [GETPRICE.COM/AU](https://getprice.com/au)

PREMIUM LAPTOPS

APPLE MACBOOK PRO 14IN

STUNNING CREATIVE POWER

Alongside its 16-inch sibling, this is simply the world's best laptop for demanding users. The amount of power on tap via the M1 Pro or M1 Max processor is staggering, and it's backed up by a terrific screen and stunning battery life. With all the ports you could ask for, the only negatives are its relative bulk (especially compared to the Air) and high price.

From \$2,999, apple.com/au



ALTERNATIVES

DELL XPS 15 (2022)

Surely the best Windows all-rounder money can buy, with a slim design, gaming potential and plenty of power courtesy of Intel's 12th gen Core processors.

From \$2,924, dell.com

APPLE MACBOOK AIR

The MacBook Air M2 looks like a standout to us. A powerful and quiet ultraportable with an unbeatable battery life and a great pro display.

From \$1,899, apple.com/au

ASUS VIVOBOOK PRO 16X OLED

Top notch components, an amazing OLED display and leading battery life make this an exceptional Ultrabook for creatives.

\$2,299, asus.com/au

BUSINESS LAPTOPS

DELL XPS 13 PLUS

DELL'S NEW FLAGSHIP IS A STUNNER.

The updated XPS's headline features include a notably more powerful Intel 12th generation CPU, an optional OLED touchscreen panel, an integrated glass touchpad, palmrest and a neat capacitive touch bar. A powerful luxury laptop that's up there with the very best.

\$2,339, dell.com/en-au



ALTERNATIVES

LENOVO THINKPAD Z13 GEN 1

The perfect laptop for executives, with amazing battery life backed up by strong performance. A great debut for the Ryzen Pro 6850U.

\$2,929, lenovo.com/au

LENOVO THINKPAD X1 CARBON GEN 9

A business laptop with heaps of luxury, especially if you choose the top-end model with a 3,840 x 2,400 screen.

\$2,959, lenovo.com/au

DYNABOOK PORTEGE X30W-J

This 2-in-1 convertible is ludicrously light and slim, and offers a quiet keyboard, great performance, ports a-plenty, and a nice touchscreen.

\$2,083, anz.dynabook.com

GAMING LAPTOPS

NEW ENTRY

RAZER BLADE 15 (2022)

A SUPERBLY ENGINEERED GAMING POWERHOUSE

Right now there's just not another better gaming unit around. It's miraculously thin considering the beefy components inside, thermals are very well managed and the Razer Blade 15 manages to extract very good performance from its GPU without undue throttling. There are many variants and options, and they're an outstanding buy if you find one on sale. ~\$4,499, razer.com/au-en



ALTERNATIVES

LENOVO LEGION 5I PRO (16IN)

A great-value gaming laptop that's extracts the most from its powerful components. We really love the keyboard, too.

From \$2,649, lenovo.com/au

ASUS ROG STRIX SCAR 17 SE

An extremely high performance esports gaming laptop that was a good 15 percent faster than any other laptop we had on file.

\$5,799, asus.com/au

ASUS ROG ZEPHYRUS G14 (2022)

A powerful gaming ultraportable that lights up for both work and play. The AMD CPU and GPU balance well.

\$3,499, asus.com/au

EVERYDAY LAPTOPS

MICROSOFT SURFACE LAPTOP GO 2

FEELS LIKE A PREMIUM ULTRABOOK, BUT COSTS HALF AS MUCH.

A solid update to an amazingly good value mid range ultrabook. It runs cool and quiet, and has up to 30 percent better battery performance than its predecessor, all in a beautifully designed form with a superb keyboard and trackpad.

From \$1,099, microsoft.com/en-au



ALTERNATIVES

LENOVO YOGA SLIM 7 PRO

A powerful, well designed ultrabook with an awesome OLED screen for a good price. The original 2-in-1 bends into a new clamshell shape.

\$1,699, lenovo.com

HP ENVY 13 (2020)

It may not be the world's most exciting laptop, but the Envy 13 is a compact and forward-looking machine that's ideally suited to these homeworking times.

\$1,699, store.hp.com

HP PAVILION 14

This slim 1.4kg laptop is a superb choice for anyone on a tighter budget. It struck an excellent balance between speed and battery life – around ten hours.

From \$999, store.hp.com

CHROMEBOOKS

ASUS CHROMEBOOK FLIP CX5

The best big-screen Chromebook we've tried, with a bright 15.6in Full HD display with decent black levels and surprisingly rich colours. And it's a good specification for the price, with a Core i3 processor, 8GB of RAM and a 128GB SSD in tow. \$1,213, asus.com/au



LENOVO FLEX 3 CHROMEBOOK

This isn't the most powerful Chromebook, but students need reliability more than power – especially if they're only really using it for homework, note-taking, and perhaps the occasional bit of Netflix. \$484, lenovo.com/au

LENOVO IDEAPAD DUET

The Chromebook answer to Microsoft's Surface tablets, this is a seriously versatile device – albeit not the speediest (although we never found it to be painfully slow while carrying out everyday tasks). For this price, and with a 16hrs 14mins battery life, the Duet is a great choice. \$899, lenovo.com/au

TABLETS

APPLE IPAD PRO 12.9IN (2021)

M1 POWER IN YOUR PALM FROM \$1,199
Blurring the boundaries between laptop and tablet like never before, the M1-powered iPad Pro 12.9-inch is every bit as fast as we expected – and the "Liquid Retina XDR" display is simply phenomenal. From \$1,649, apple.com/au



APPLE IPAD (2021)

Apple performs its usual trick of upgrading just enough to keep the iPad affordable and up to date, without treading on the toes of its more luxurious stablemates. The big bezels remain, but a faster processor, more storage and improved camera all add to its appeal. From \$499, apple.com/au

SAMSUNG GALAXY S8 ULTRA

The Samsung Galaxy Tab S8 Ultra is an enormously impressive tablet. It makes a fantastic first impression with its gorgeous screen and thin body, and is great for multitasking. If you want the best Android tablet going, there's no question at all that this is it. \$1,499, samsung.com

EVERYDAY MONITORS

PHILIPS BRILLIANCE 346P1C

34-INCH CURVED DOCKING, No-one can match Philips for its mix of quality, features and value, especially when it comes to curved MVA monitors such as this. It's great for casual gamers and busy workers alike, with USB-C docking and superb colours. \$849, philips.com.au



EIZO FLEXSCAN EV2480

Businesses in particular will love this easy-to-roll-out 24-inch monitor, with its exceptional five-year warranty, USB-C port for easy docking, excellent viewing angles and delicious whites. And it has the best OSD around. \$699, eizoglobal.com

DAHUA DHI-LM28-F420

This affordable 4K monitor might not offer much in the way of menu options or high-end gaming perks, but it offers impressive resolution and colour for all pay. \$269, dahuaaustralia.com.au

PROFESSIONAL MONITORS

EIZO COLOREDGE CG319X

CREATIVE MASTERCLASS, \$7,119
As the price indicates, this monitor is for heavyweight creatives who demand the best in every discipline: HDR video editing, print layouts, professional photography and more besides. With superb coverage and accuracy across all spaces, plus a built-in calibrator, it justifies the investment. \$7,119, eizoglobal.com



BENQ PD2725U

By no means a cheap 4K 27-inch monitor – unless you compare it to the Eizos – but it marries all-round quality with ease of use thanks to a puck that allows you to quickly move between settings. You can even daisy chain a second Thunderbolt 3 monitor for a monster setup. \$1,799, benq.com

EIZO COLOREDGE CG279X

Designers who need to work across different disciplines will love how easy it is to switch between the Adobe RGB, DCI-P3 and sRGB colour spaces using the Eizo's fantastic OSD. It's certainly not cheap for a 27-inch 1440p monitor, but it's packed with quality. \$3,095, eizoglobal.com

HOME OFFICE PRINTERS

EPSON ECOTANK ET-1810

BACK TO BASICS
There's nothing flashy about this bottle-fed inkjet: it's there to churn out pages at a decent rate (8.8ppm for black text) for a low price. It comes with enough ink for 4,500 black pages and 7,500 in colour, but refills are cheap. Just don't expect features like duplex printing. \$299, epson.com.au



CANON PIXMA PRO-200

Yes, it's big and it's expensive – and you don't even get a scanner – but the Pixma Pro-200 means that you can print A3 photos with glorious colour accuracy and detail. If you're after an office workhorse, look elsewhere, but it can't be beaten on photo quality. \$809, canon.com

BROTHER MFC-J4540DWXL

Home workers will love this inkjet all-in-one. It combines an incredible range of features with all the connectivity you need and extreme ease of use. Output quality is fine, it offers the best cloud support around and the high-capacity ink pack could keep you going for years. \$420, brother.com

WIRELESS ROUTERS

SYNOLOGY WRX560

THE PERFECT ROUTER

Its quad-band Wi-Fi smashed-through our performance tests managing 791.5Mbps up close, 375Mbps two rooms away and an amazing 216.7Mbps down the garden.. \$599, synology.com/au



ASUS RT-AX82U

The impressive RT-AX82U punches well above its price when it comes to raw Wi-Fi 6 performance and includes a fistful of useful features – such as an ultrafast 20Gbits/sec USB connector, outbound VPN support, site-blocking and a host of parental controls. \$320, asus.com/au

SYNOLOGY RT6600AX

Fast, well-featured and secure, the Synology RT6600ax is as perfect for professionals and prosumers as much as it is for SMBs and smart homes, and the app and interface are a delight to use for both experts and plug and play users alike. \$599, synology.com/au

MESH WI-FI

TP-LINK DECO XE75 AXE5400

SUPERB WI-FI 6 MESH

Offers the simplicity, speed and range most people want from a mesh system. While it's expensive, nothing else offers the plug-and-play simplicity, speed and range afforded by the Deco XE75. \$999, tp-link.com



TP-LINK DECO X20

The price below isn't a typo: the Deco X20 makes it possible to upgrade your home network to Wi-Fi 6 – as well as manage its security – for the price of some standalone routers. The included HomeCare functions are also a great bonus for households where lots of devices are fighting over the bandwidth. \$199, tp-link.com

ASUS ZENWIFI AX


The interface and design certainly aren't the prettiest (the nodes look like mini space heaters, while the charcoal version is straight out of the 1980s), but if you're after strong performance, great features and plenty of expansion potential, the ZenWiFi AX is a compelling prospect. \$949, asus.com/au

NAS SERVERS

QNAP TS-130

SIMPLE HOME NAS

An above average CPU and RAM combo, solid features and affordable price make the QNAP TS-130 worth considering. Alternatively there's the two bay TS230 which features the same CPU, with 2GB of RAM, dual drive bays, and an extra USB port on the front for a modest \$50 price increase. \$250, qnap.com



ASUSTOR NIMBUSTOR AS5304T

This is a hugely powerful four-bay NAS that will be overkill for most people – not just due to its speed but also its slightly intimidating user interface – but techies and gamers should investigate the Nimbustor AS5304T for both its speed and flexibility. \$799, asustor.com

SYNOLOGY DISKSTATION DS1522+

At \$1,199 it's not cheap but factoring in the software means it's great value. Though some internal silicon may lack headroom, this is still a well-featured, expandable NAS with fantastic software. The best part is the DSM software. \$1,199, synology.com

SECURITY SOFTWARE

F-SECURE SAFE

F-Secure concentrates on the basics: faultless protection with minimal impact on your system's performance – for a low price. 1-3 PCs, 1yr, \$99.99, f-secure.com



ALTERNATIVES

NORTON 360

Consistently on the podium in our group tests, and loaded with an abundance of genuinely useful features, including a VPN and generous cloud backup. 1 PC, 1yr, \$79.99, nortonlifelockpartner.com

AVAST FREE ANTIVIRUS

For the second year running, Avast is our top pick for free antivirus software – it's simply lighter on its feet than Microsoft Windows Security, whilst offering fewer false positives. Free, avast.com

VPNs

NORDVPN

NordVPN provides consistent and fast speeds, serious security, great support for video-streaming services and some cost-effective subscription rates. \$59 per year, nordvpn.com



ALTERNATIVES

PROTONVPN

ProtonVPN provided one of the best free offerings of all the VPNs in our group test, including unlimited data, but upgrade to benefit from even faster speeds and many more options. Free, protonvpn.com

SURFSHARK

A strong rival to NordVPN, especially if you're willing to commit to its two-year contract. It's fast, cheap and a fine choice for people who like to switch to US streaming services. \$66.34 for one year, surfshark.com

PASSWORD MANAGERS

BITWARDEN

Bitwarden has a huge advantage: it's free. It isn't as slick as some paid-for rivals, but it can sync passwords across all devices for no extra charge. Free, bitwarden.com



ALTERNATIVES

DASHLANE

A manager that's ideal for beginners, and it even builds in an unlimited (if basic) VPN service. Note you may prefer to buy the Family plan (\$60 per year) as this extends the service to six people. \$55.99 per year (Premium), dashlane.com

1PASSWORD

1Password is targeted at technically minded users who are looking for the last word in security. It even offers a Travel Mode that may ease your mind if surrendering your phone to customs officials. \$54.99 per year (individual), 1password.com

THE LIST

The best power supplies for PC gaming

While the rumoured 600W+ power requirement for the RTX 4090 didn't quite pan out, Nvidia does recommend at least an 850W power supply for this generation's top card, up from 750W for the 3090. So there are probably quite a few people out there looking to upgrade their rig. Regardless of whether you're buying a new system or upgrading though, it's always a good idea to have a good PSU at the heart of your system. These are five of the PSUs we recommend for gamers.



1

CORSAIR RM750X

www.corsair.com | \$199

While we're all getting carried away at present with 850 and 1,000W PSUs, the truth is that 750W will still be more than enough for most people. The RMx line is among the best in the PSU market, combining high performance and reasonable prices. This model uses a magnetic levitation fan for quiet operation, plus high-quality Japanese caps, and has a long 10-year warranty.



2

SEASONIC PRIME TITANIUM TX-1000

www.seasonic.com | \$529

If you're considering a high-end GPU upgrade in the next couple of years then it makes sense to go for at least 850W, if not 1,000W for proper future-proofing. Seasonic's base platform is used by many other manufacturers including Corsair. This PSU is crazy efficient and quiet, thanks to its fluid dynamic bearing fan, and also comes with a generous 12-year warranty.



3

BEQUIET! PURE POWER 11 FM 550W

www.bequiet.com | \$170

If you're running an average system without the need for stupid power requirements, this PSU from BeQuiet! is the one to go for. It's 80 PLUS Gold and Cybenetics Gold-certified in efficiency, plus it has a Cybenetics A noise rating, meaning it's ultra-quiet. But perhaps best of all is that it's priced at under \$200, which is excellent value for such a high-quality PSU.



4

CORSAIR AX1600i

www.corsair.com | \$699

If you're going to be pushing your PC by filling a computer with the most power-hungry components, and then driving them hard via overclocking, then this is the PSU for you. As well as its ridiculous 1600W capacity, it's highly efficient, thanks to GaN MOSFETs, which can offer up to 99 percent efficiency. It is remarkably quiet too, thanks to its high-quality FDB fan.



5

SILVERSTONE SX1000 PLATINUM

www.silverstonetek.com | \$369

The best small-form PSU on this list doesn't stop it from having every connector a gaming PC builder needs, at least for the next few years. It's completely modular, which is handy for compact builds as you can leave out cables to save space. It can be pushed hard too – we monitored 1500W capacity in our testing. The double ball-bearing fan also keeps it cool.

HEAD TO HEAD

Google Drive vs. WeTransfer

Which side are you on for file sharing?

■ GOOGLE DRIVE | \$free | GOOGLE ONE 100GB | \$2 per month to 30TB \$300/month | drive.google.com
 ■ WETRANSFER | \$free | WETRANSFER PRO US\$10 per month | WETRANSFER PREMIUM US\$19 per month | www.wetransfer.com

Gone are the days of sending a USB stick containing work, photographs, or other important documents over to your colleagues via snail mail or carrier pigeon. If you are clinging on to the old ways, then fair enough, but times have changed and there are now much quicker and more efficient ways of sharing your files across the internet.

Thankfully, using the cloud is a much more effective and secure method of sharing data. There are plenty of cloud services out there, but two of our favourite methods are Google Drive and WeTransfer – the two titans of the data-sharing world. Although both platforms differ in many ways, at their core, they both allow you to send and receive files for free. Paid options are also available for both services and we will touch on the finer details of these later.

Both of these services are fairly recent inventions. WeTransfer was founded in 2009 by Bas Beerens and Ronald “Nalden” Hans in Amsterdam, Netherlands. The pair simply wanted to solve a problem they were experiencing at the time, which was sending large intensive design files across the internet. At its launch, WeTransfer had humble beginnings – its only function was to send files of up to 2GB in size for free.

But just three months after it

launched, WeTransfer was being used in 187 countries around the world, so it's fair to say that lots of other people had also been experiencing the same problem. Since then, the WeTransfer platform has expanded into a multitude of services, including WePresent, Paper, Teams, Portals, Collect, and Paste.

Google Drive was created in 2012 and has a slightly different identification as a file storage and cloud platform. Offering a lot more features than WeTransfer, it aimed to be the all-in-one platform for sending, receiving, and managing files, either for personal or business use. Of course, being a Google product, it integrates nicely with the rest of the Google Suite, creating an almost seamless ecosystem.

That said, if you are already a part of this suite, is WeTransfer a better bet for efficient file sharing, or is it worth sticking with Google Drive? These two services have similar core abilities but there are a few key differences. There is more than meets the eye with these applications so we're going to put the free offerings head-to-head to see which is more worthy of your precious uploading time.


On paper, you might be happy with the features that a service offers and the way they've been laid out. However, if you then use that service and it's a

clunky, unorganised mess of an application, chances are you aren't going to want to use it. That's presumably what the team at WeTransfer had in mind when they designed its UI. Based upon founder Bas Beerens's creative studios, OY Communications, from the start, WeTransfer featured a clean interface and seamless navigation, creating a smooth experience for the users. Those characteristics are still prevalent today, and we found the WeTransfer site to be well-designed and easy to use.


The main part of the landing page includes a box where you can upload your files and send them. There are further links and information in a bar along the top right corner of the page and other tools in a drop-down icon from the top left. The majority of the page includes an advert that is often live or interactive and is seamlessly designed with WeTransfer in mind. Although not everyone wants to look at an advert, it is located in a place where it's not too intrusive and we don't mind it being there given the free service.

Ad-free service

However, when using Google Drive, we aren't greeted by any ads, which is still the preferable option. This service follows the same design language that Google has implemented across the board over the past decade – a



Google Drive provides an effective all-in-one home for your file management needs.



WeTransfer is a slick and efficient file sharing tool, if you don't mind staring at an advert.

minimalist style that uses the colours from Google's branding. We're big fans – the theme ties Google's products together and just works.

When you arrive on the Google Drive landing page, you're greeted by folders and your most recent files displayed in square icons, making it easy to distinguish recent or frequently opened files. On the left-hand side, there's a list of functions, including your drive, shared drives, items shared with you, recent, starred, and the bin. Above that is the 'new' icon where you can create spaces to upload and share files with others. A large search bar across the top makes it easy to find files, especially as it includes a fine-tuning button to narrow down your search. The top right features your Google account and access to other applications within the Google suite, which is handy.

Despite its solid structure, those who are new to the Google universe may find it takes a while to get used to. On the other hand, WeTransfer makes the whole process extremely simple, while also being easy on the eye, so we have to hand the design element to WeTransfer here.

When it comes to overall features, Google Drive is a much more well-rounded product. It offers you a place in which you can create your own shared file folders, and where other users can then contribute their own files and documents. If you create a folder, you can then choose who else has access to it by inviting people via an email link.

Google Drive is free, but there is a total cap of 15GB, which restricts its usefulness somewhat. The 15GB cap is

presumably there to entice you to upgrade to Google One, where you can pay monthly for cloud storage space across the Google Suite all the way from \$2 a month for 20GB to \$300 a month for 30TB. There are plenty more plans available between these smallest and largest monthly payments, giving you lots of options as you increase the amount of storage you need.

With WeTransfer, you can either get a specific link that you can share or you can send files directly from the WeTransfer homepage. The link expires after one week, during which time it allows the recipient to open and download the files you have sent. Alternatively, you can send files via email from the WeTransfer homepage. You can track your transfer history via the transfers tab at the top right of the page, allowing you to keep an eye on the files you have sent.

Just as Google Drive has a range of pricing options, so WeTransfer has some paid options. The Pro package costs US\$10 a month and lets you send files of up to 200GB capacity at a time and gives you 1TB of cloud storage. The Premium package is US\$19 per month and you get an unlimited amount of cloud storage. You also get to send huge files with no limit on the capacity. In our opinion, that makes WeTransfer better value for money than Google Drive. For just under US\$20 per month, you have unlimited cloud storage and the ability to send huge files across the internet. However, it's evident that Google Drive's free package has many more features than WeTransfer's does.

When testing, we were impressed with the fluidity, effectiveness, and efficiency of both platforms, but more

so with WeTransfer. Although WeTransfer offers far fewer features compared with Google Drive, it does have a simple UI that eliminates any clutter, and while you have to put up with adverts on the homepage, they're well-designed and fit neatly with the interface.

If ease of use appeals more to you, then WeTransfer could be a better bet. It prioritizes user-to-user file sharing, whereas Google Drive is more focused on creating a space where you can share files. However, as a free package, Google Drive offers everything WeTransfer does and much more besides.

Although both services are great tools for sharing data, WeTransfer works as a quicker alternative for single files, while Google Drive has the benefit of allowing you to organise your files better. It also makes more sense if you are using other Google productivity services, such as Gmail and Google Docs, as Google Drive can integrate these into your daily life. The well-roundedness of Drive makes it our file-sharing service winner.

Google Drive Well-known UI; easy file management; safe and secure. 15GB cap; encourages use of other Google services.

Sam Lewis

★★★★★

WeTransfer Simple well-designed UI; efficient solution. 2GB file size cap is poor and it features adverts.

Sam Lewis

★★★★☆

TECHNOTES

Gadgets

Lovely little things



1

GMC HUMMER EV AWD SUPER EBIKE

The perfect ride for militant-environmentalists.

US\$3,999 | reconpowerbikes.com/hummer-ev-bike

We have to say, we didn't expect to see the all-American diesel-guzzling Hummer ever go all electric, but this military-made offroader has converted its 2023 offering to run exclusively on electric power and you can even get one with its very own eBike. Sure it's not quite the Cybertruck and ATV combo Tesla showed off, but they're a lot more available, with the all-wheel-drive ebike launching in January and the car set to be released later in 2023. The GMC Hummer EV AWD Super Ebike's 17.5Ah battery offers 60km of range to get anywhere this electric ute can't quite reach.

2

SEALSHIELD CLEANWIPE PRO WATERPROOF KEYBOARD

A keyboard that's clean enough to eat off.

US\$150 | sealshield.com

Sure, your mechanical keyboard might have N-key rollover, fancy RGB backlighting and even customisable keys, but all these features won't help it from becoming utterly useless when you spill your banana smoothie on it. Thankfully the Cleanwipe Pro Waterproof Keyboard is crumb-resistant, sauce-proof, splash-safe, sneeze-repellant and entirely washable from top to toe. You can even chuck this IPx8-certified keyboard in the dishwasher...

and if that doesn't work, bathe it in bleach, according to the label. We won't ask what you did to the keyboard for it to require this level of cleaning, but we can guarantee whatever Cherry's you're tapping away on now won't be so durable.





3 RAZER HAMMERHEAD HYPERSPEED

Gaming ear pods that drive home communication.
\$285.95 | razer.com/en-au

Razer's new Hammerhead Hyperspeed headphones are designed to combine the convenience of wireless Bluetooth earbuds with the reliability of a dongle connected headset for gaming. In addition to offering Bluetooth 5.2, the Hammerhead Hyperspeed comes with a compact dongle to connect to a PlayStation 5 Dual Sense controller (or any USB-C port) to allow you uninterrupted comms over WiFi 2.4GHz. These headphones also include active noise cancelling and environmental noise cancelling modes to help you, and whoever you're talking to, hear what's important. The headphones offer 4.5 hours of charge (ANC on) and carry an additional four charges in the case.

4 SEGWAY NINEBOT MAX 2

An e-scooter to put you on cloud Ninebot.
\$1,199 | au.segway.com

Sure, it might not be the coolest looking electric vehicle we've seen, but the Segway Ninebot Max 2 is a comfortable and compact mode of transport that's going to appeal to a lot of people. With a 65km range it'll get you from A to B for the better part of a week before you'll need to recharge, and with a top speed of 25km/h it'll get you where you need to be in no time. 10-inch, puncture proof, pneumatic tyres provide heaps of cushioning and the powerful motor can push up inclines of up to 20 percent.



6 DJI MAVIC 3 CLASSIC

The top cinematography drones at a more affordable price.
\$2,399 | store.dji.com/au

The new Mavic 3 Classic from DJI offers the same updated camera, a new controller and rejigged components as the Mavic 3, but lands at a price point that's \$500 less. This new device keeps the 4/3 CMOS sensor and precision L2D-20c lens for shooting 5.1/K HDR video footage in a wide range of lighting (or night-time) conditions. It's also able to stay in the air for up to 45 minutes and offers high quality 4x digital zoom to get that little bit closer to what you want to shoot. It's even got improved automatic object-tracking making it easier to keep what you want in shot. won't be so durable.



LAPTOP

Razer Blade 15 (2022)

Is the Blade 15 still the best tool for gamers on the go, or has the device dulled?


■ \$4,499 | razer.com/au-en

The Razer Blade 15 really isn't just one device. You have a choice between four processors (spanning three generations), five graphics cards, four gaming screens, two RAM configurations, and two storage configurations – so the Blade 15 is really an entire range of laptops. This versatility, particularly in a screen type, is part of the reason why it's frequently in the top spots of the best gaming roundups, but there's another factor to consider: it's extremely compact.

The Razer Blade 15 is as thin as the LG Gram 16-inch 2-in-1, and 100g lighter than a MacBook Air 16... even when configured with a Nvidia GeForce RTX 3080 Ti. Normally squeezing an enormous GPU like this into an Ultrabook form means you have to sacrifice a lot in performance. And on paper the 110W max power draw of the Blade 15's 3080 Ti GPU seems a lot lower than the 165-175W power draws that bigger gaming rigs use to power this card, but the reality is that the performance gap is not as big as you'd expect here.

The synthetic graphical benchmarks do show a performance lag of up to 25 percent against some units, but the Blade 15 performed en-par in almost all GPU tests with the 165W 3080 Ti on a Asus ROG Zephyrus Duo 16. This means playable (40fps) frame rates on Cyberpunk 2077 using 1440p Ray Tracing Ultra settings and over 110 fps averages on less demanding games at 1080p Ultra.

When it comes to the CPU, the model we tested came with an Intel Core i7-12800H. You can configure it with an i9-12900H for those wanting a bit more power, but it won't make too much of a difference unless you



■ W11 Home; 15.6-inch 240Hz panel at 2560 x 1440 pixel resolution; Intel Core i7-12800H CPU; Nvidia GeForce RTX 3080 Ti GPU; 32GB RAM; 1TB PCIe SSD; 81Wh battery (4h 54min 1080p movie playback); 35.5 x 23.5 x 1.7 cm; 2.1kg.

plan to do a lot of CPU heavy content streaming or media encoding. Razer isn't afraid to sell 11th and even 10th Gen Intel CPUs alongside the current range for a marked down price...but 12th Gen was a big step forward for Intel, so it's a bit of a false economy buying one of these older models, even if it is thousands less.

That said, the one area that is still a major sticking point for the Blade 15 is the price. The RRP's for the 12th Gen units range from \$4,549 for a 3060 model to \$7,199 for an i9, 3080 Ti model with a 4K 144Hz panel. At this price the devices aren't overly competitive, but you can get them on sale if you keep an eye out. The Blade

15 we tested (QHD 240Hz, i7, 32GB RAM, 3080 Ti, 1TB) had an RRP of \$6,699, but was available for just \$4,499 if you shopped around a bit.

Battery life was pretty average for this kind of unit, lasting between four and five hours for general work tasks or 1080p movie playback. The Trackpad and keyboard are great and since you can choose from monitors including: a FullHD, 360Hz panel; a QHD, 240Hz, DCI-P3 G-Sync monitor; a 4K, 144Hz, DCI-P3 screen; or a 240Hz QHD OLED display, you can get the perfect visuals for whatever you need to do.

While the cooling on this unit is impressive and not abnormally loud you should expect 100-degree CPU spikes and a bit of thermal leakage on the top keyboard surround, which can be hot to touch. If you can find these units on sale then we can't think of another better gaming unit around.

On sale, the 2022 Blade 15 is the best performance-per-kg gaming laptop on the market.

Joel Burgess

★★★★★

	Razer Blade 15 (2022)	Asus ROG Zephyrus Duo 16
PCMark 10 - Overall (score)	7,126	7,114
Cinebench R23 - CPU (Score)	10,796	14,275
Max CPU Temperature (C)	100	85
1080p movie playback (HH:MM)	04:54	02:50
3DMark Time Spy (score)	9,950	10,180
F1 2021 - 1080p Ultra (fps Av.)	116	117
CrystalDiskMark Read (MB/s)	6,907	6,552
CrystalDiskMark Write (MB/s)	4,951	4,399

LAPTOP

Dell Inspiron 16 Plus 7620

Is Dell's new Inspiron Plus a budget MacBook Pro killer or just a plus-size LG Gram 16?

■ \$3,397.90 | dell.com/en-au

There's often many ways to frame a laptop, depending on what you compare it to. From one perspective, the Dell Inspiron 16 Plus 7620 is a heavy 2,188g alternative to the 1,480g LG Gram 2-in-1 16. From another it's a much cheaper Apple 16-inch MacBook Pro and from another again it's an expensive and late product that's comparable to an MSI Crosshair 15. Each of these comparisons have trade-offs and caveats and one is undoubtedly more flattering than the others, but all of them are true and, depending on what you're looking for, will dramatically transform how good you think the Dell Inspiron 16 7620 is.

The Inspiron 16 Plus is a large laptop aimed at professional ultrabook users if you go by the marketing material. The 16-inch 3072 x 1920 display is more detailed than your average 15.6-inch PC display, lining up pretty closely with the Apple MacBook Pro 16 resolution, but it's only got a 300 nit TN panel with no colour certifications so in many ways it's perhaps incomparable to the 1,000 nit 120Hz DCI-P3 Apple panel.

There's a couple of configurations of the Dell Inspiron 16 Plus, but they're all built around an Intel Core i7-12700H CPU. This is a powerful processor, commonly used on gaming laptops and other bigger work laptops with beefier cooling architecture than what you'll see on an Ultrabook. You can then pair that

"The 16-inch 3072 x 1920 display is more detailed than your average 15.6-inch PC display, lining up pretty closely with the Apple MacBook Pro 16 resolution."

with either a Nvidia GeForce RTX 3050 or an RTX 3060 GPU; 16 or 32GB of RAM; and choose from 512GB or 1TB of PCIe SSD storage. The entry device here lands at \$2,598.20, while maxing out on all these options will set you back \$3,397.90. This isn't a bad price for what you get, seeing as it'll perform similarly to a \$4K 16-inch MacBook Pro, but we've seen similar specced MSI

Crosshair 15 laptops on sale for \$2,599.

The performance on the Inspiron 16 Plus is really solid, offering between

36 and 220 percent improvements over something like the LG Gram 2-in-1 16 on various general work benchmarks. This lines up where we'd expect against more powerful work and gaming laptops and when paired with a RTX 3060 and 32GB of RAM the gaming and graphical capabilities are respectable. The 60Hz display isn't ideal for gamers, but you can play modern titles at 30+ fps on QuadHD Ultra presets thanks to the discrete GPU, which will make it a solid option for graphically intensive workloads.

Battery life with this one is a bit of a mixed bag since you'll get over 10 hours from optimised processes like 1080p movie playback, but just 3 hours and 3 minutes of active work time thanks to Nvidia Advanced Optimus GPU switching.

The keyboard and trackpad are nice and it's got a usable webcam and mic/speaker array. The unit does have a few hot spots around the rear base and on the keyboard surround that are uncomfortable when you push the machine, but the processor peak temperatures were fine, peaking at 96-degrees.

A powerful MacBook Pro 16 alternative if you don't need a colour accurate screen.

Joel Burgess

★★★★☆



■ W11 Home; 16-inch TN WVA display at 3072x1920 pixel resolution; Intel Core i7-12700H CPU; Nvidia GeForce RTX 3060 GPU; 32GB RAM; 1TB PCIe SSD; 84Wh battery (10h 12min 1080p movie playback); 35.7x25.2x1.7cm; 2.2kg.

	Dell Inspiron 16 Plus 7620	LG Gram 2-in-1 16
PCMark 10 - Overall (score)	6,995	5,125
Geekbench 5 - Multi-core (score)	12,599	7,792
Cinebench R23 - CPU (multi-threaded)	15,079	6,820
3DMark - Time Spy (score)	7,118	1,631
Battery life - 1080p video playback (h:min)	10h12min	13h22min
Peak CPU temperature (°C)	96	100
CrystalDiskMark Read (MB/s)	5,274	6,951
CrystalDiskMark Write (MB/s)	4,938	4,350

ULTRAPORTABLE LAPTOP

HP Envy X360 2-in-1 Laptop 13

Does HP's latest convertible have enough business perks to close the deal?

■ \$2,699 | hp.com/au-en

The HP Envy x360 13 2-in-1 is a little late to the 12th Gen party, launching a good four months after market leaders like the Dell XPS 13 – but HP has used this extra time well, tweaking this 2-in-1 into a product that consumers will love. The top-spec version of this 13.3-inch touchscreen convertible utilises a 2880 by 1800 pixel OLED panel that looks exceptionally rich. That panel also features a 120Hz refresh rate, which enhances the responsiveness of the touchscreen when drawing or writing with a stylus, and the full DCI-P3 colour capabilities make it a solid option for creatives.

When it comes to appearances, the HP Envy x360 13 looks like a classic business 2-in-1. It's got a boxy profile with rounded edges and a clean silver metallic case. The classic 2-in-1 hinge design is simple and allows full 360-degree rotation, which makes the accompanying rechargeable USB-C stylus comfortable to use.

The RRP for this top spec unit is \$2,699, which gets you the OLED display, an Intel Core i7-1250P CPU, 16GB of RAM and a 512GB PCIe SSD, but we did see this device on sale for just \$2,159 on HP's local site at the time of writing. If your budget can't stretch to this there's also an i5 version with 16GB of RAM and a 16-by-10 1920 x 1200 pixel IPS display for \$1,999; or a Ryzen 5 model with 8GB RAM, a FHD IPS display and a 256GB PCIe SSD for \$1,499 (\$999 on sale).

The top spec Envy x360 13 2-in-1 is definitely the pick of the bunch here offering a standout screen without adding too much to the price tag. The Envy x360 is able to do this by using an Intel Core i7-1250P CPU which is a bit cheaper since it only has two performance cores – unlike the i7-1260P which has four or the i7-1280P



■ W11 Home; 13.3-inch OLED display at 2880 x 1800 pixels, 100% DCI-P3 colour, HDR 500; Intel Core i7-1250P CPU; Intel Iris Graphics; 16GB RAM; 512GB PCIe SSD; 69Wh battery (8h33min 1080p movie playback); 29.8x21.5x1.6cm; 1.34kg.

which has six. This puts it up to 16 and 19 percent behind the 1260P in CPU and GPU benchmarks respectively and closer to 30 percent behind the 1280P on both. For the average worker this isn't a big deal seeing as the new processor is still significantly faster than any comparable 11th Gen Intel processors, so it's more than capable of editing photos, running a stack of browser tabs and smashing office applications. That performance cut does, however, mean it isn't ideal for gaming, but if you like the look of this device you're more likely to prefer doing a tax return with a bag of plainly salted crisps than waste valuable time on a frivolous videogame.

The Envy x360 13 does come with Windows 11 Home and we couldn't see an option to swap that out for W11 Pro, but you can add two or three year on-site device servicing for \$199 and \$249 for anyone concerned about it failing as a primary work device.

Whoever designed the laptop didn't account for the extra thickness of an Organic LED lighting array over a full touchscreen panel, since the screen protrudes a good few millimetres above the metal clamshell casing. It didn't create any immediate concerns, but we do worry a little about longevity and resistance since the screen is less protected than your average device. There was also no fingerprint reader – an odd omission on such a business focussed device – but there is at least Windows Hello facial recognition for speedy sign-ins. The Gen 4 PCIe SSD used offers 6,800 and 4,400MB/s respective read and write speeds and battery life just ticks over the eight hour mark for both work and 1080p movie playback.

	HP Envy X360 2-in-1 Laptop 13	Dell XPS 13 Plus
PCMark 10 - Overall (score)	4,879	5,125
Geekbench 5 - Multi-core (score)	7,192	8,594
Cinebench R23 - CPU (multi-threaded)	7,667	8,161
3DMark - Time Spy (score)	1,475	1,853
Battery life - 1080p video playback (h:min)	8h33min	6h6min
Peak CPU temperature (°C)	96	100
CrystalDiskMark Read (MB/s)	6,817	7,082
CrystalDiskMark Write (MB/s)	4,467	5,095

A reasonably priced 13-inch work convertible with an OLED display.
Joel Burgess

★★★★☆

Intel NUC 13 Extreme Raptor Canyon

Is the new benchmark for compact gaming PCs something all PC builders should consider?

■ \$2,899 | intel.com.au

Sitting somewhere in between a DIY build and a ready-to-go prefabricated gaming rig, Intel's Next Unit of Computing (NUC) Extreme devices are designed to showcase what's possible with the latest gaming-PC components. The NUC 13 Extreme Raptor Canyon is available with an Intel 13th generation i5, i7 or i9 processor and includes a motherboard, case and a 750W 80+ Gold SFX internal power supply – which is maybe 50-percent of what you need to fire up a computer. The remaining components are left for you to buy and install yourself, based on exactly how much power you need.

The i9 units were on pre-sale at the time of writing for \$2,899, but you'll have to add up the cost of RAM, storage, an OS, and a discrete GPU to get a full picture on what it'll cost before even considering a screen and peripherals. Seeing as you can get really solid all-inclusive gaming PCs for that price, the NUC 13 Extreme is not attempting to appeal to the DIY-builders on a budget, it's more aimed at PC performance enthusiasts that are willing to pay top dollar for the most compact high-end PC possible.

If you do the maths on the i9-13900K (\$1,099), a mini ITX motherboard (\$300-\$600), custom CPU fan and heatsink (~\$200), power supply (\$150), and a case (~\$150), there's about \$700 left over that you're paying Intel for the engineering devoted to ensuring that this is the most compact high-performance gaming rig around. There's undoubtedly a fair amount of work in guaranteeing a unit will efficiently move air through it and that it'll work with the latest graphics cards in a



■ Custom Intel NUC 13 Extreme Compute Element and motherboard (30.5x13.7x4.7cm), Unlocked 13th Gen i5, i7 or i9 CPU; Dual channel DDR5 SODIMMs (max 64GB); 3 x m.2 key m PCIe SSD (Gen 4); PCIe x16 Gen5 slot; 750W power supply included; Intel Wi-Fi 6E AX1690i; 33.7x31.8x12.9cm chassis.

compact vertical mounting orientation.

The Intel NUC 13 Extreme Raptor Canyon is also extremely efficient at saving space; it's still about twice the size of current gaming consoles, but that equates to less than half the volume of your average mini-ITX case, without any compromise in performance, so it's not something to be scoffed at. As impressive as all these features are, there's still one glaring caveat: it's indistinguishable in appearance from a work PC that you could pick up for \$800 from Officeworks.

Honestly, if you care enough to pay thousands extra to make your PC fractionally smaller, the thought of then plastering that PC with gaudy gamer skulls makes no sense at all,

but we can't help but feel conflicted about such powerful components being hidden away in a plain black box. If you feel the same, there's no shortage of glass-lined RGB-filled mini-itx towers to show off your pride and joy, but the humble NUC 13 Extreme Raptor Canyon is a rare find for anyone who would rather let their computer's performance make the statements.

The unit we were sent to review featured an i9-13900K CPU, 32GB of RAM and a Nvidia GeForce RTX 3080 Ti (350W), so unsurprisingly it's one of the best performing devices we have on file. It managed to outpace an i9-12900KF by up to 60-percent in some thread-heavy workloads and even kept up with the RTX 3090 that was installed in that Alienware Aurora R13 device. This means 60fps averages on 4K Ultra Ray Tracing Cyberpunk 2077 and 120fps minimums on 1080p Ultra game and since you're playing with unlocked components you can push those figures further with customisation... so you'll want to factor in the cost of a fancy screen.

The perfect modular PC for discrete users
with deep pockets.
Joel Burgess

★★★★☆

	Intel NUC 13 Extreme Raptor Canyon	Alienware Aurora R13
PCMark 10 - Overall (score)	8,772	8,448
Geekbench 5 - Multi-core (score)	22,708	17,225
Max CPU Temp (C)	100	100
Cinebench R23 - multi-threaded CPU	35,382	22,029
3DMark Time Spy (score)	18,851	18,361
Sid Meier's Civilization VI - Ultra 1080p (fps Av.)	327	321
CrystalDiskMark Read (MB/s)	7,082	4,380
CrystalDiskMark Write (MB/s)	6,109	4,380

GRAPHICS CARD

MSI GeForce RTX 4080 Suprim X

The excellent RTX 4080 gets the Suprim treatment.

■ \$2,599 | www.msi.com

Hot on the heels of the monstrous RTX 4090 comes the RTX 4080. Though it should be a much more affordable option, at this early stage RTX 4080's carry prices that wouldn't have been out of place during the height of the crypto currency boom. It is what it is though. In six months, the value proposition could be totally different. Let's begin with a look at MSI's RTX 4080 Suprim X. And a real beauty it is.

MSI's Suprim X is stunning, with a brushed aluminium finish and tasteful RGB lighting.

Like all 4080's it is built around the AD103 GPU, it has 9,728 cores. It's clocked at 2,625MHz, though you can get an extra 15MHz on top of that via the MSI Center app. As is often the case though, the rated boost clocks don't mean a whole lot. In our testing, the card was able to hold a boost clock of 2,745Mhz while looping 3DMark Speedway.

The Suprim X comes with the new 12+4 pin ATX 12VHPWR connector. MSI ships a 3x8-pin adapter for use with current gen power supplies. You also get dual BIOS, with silent and gaming options.

Display connectivity consists of three DisplayPort 1.4a ports and one HDMI 2.1. Ada Lovelace cards miss out on DisplayPort 2.0/2.1.

The cooler is big and very much over spec for a 320W TDP card, but that means you get a cool and quiet running card. Shockingly, the card reached a temperature of just 57 degrees during our testing. 57! If an RTX 4090 Ti eventuates, this cooler will probably handle it.

The cooler is over spec, and so is the PCB. You get an 18-phase VRM with 50A stages. The memory is rated to run at 24Gbps, so there's some 'free' memory overclocking headroom too.

The RTX 4080 and Suprim X deliver excellent gaming performance. 4K gaming with RT is well within the capabilities of the card, especially if you enable DLSS, which is getting better all the time. It easily beats out the RTX 3080 Ti, and pulls ahead of the RTX 3090 Ti too.



■ MSI GeForce RTX 4080 Suprim X; 9,728 CUDA cores; 2,625MHz boost clock; 16GB GDDR6X 22.4Gbps memory; 716.8GB/s memory bandwidth; 3x DisplayPort 1.4a, 1x HDMI 2.1; 320W TDP, 1x 16-Pin power connector.

"4K gaming with ray tracing is well within the capabilities of the card, especially if you enable DLSS, which is getting better all the time."

One of the most impressive RTX 4080 feats is its relatively low power consumption, or more specifically its performance per watt. We recorded a maximum power draw of 317w for the Suprim X. That's lower than the RTX 3080 and RX 6900 XT, while delivering much better performance. For all the pre-release whispers of runaway RTX 40 power consumption, the RTX 4080 delivers what appears to be the best power

efficiency of any card released to date.

The RTX 4080, and particularly MSI's Suprim X flavour is stunning. The Suprim X is cool, quiet, it looks

fantastic and it has a cooler and PCB that will never be stressed by the power requirements of the AD103 GPU. But we can't ignore the price. Will it change after the imminent launch of AMD's competing Radeon RX 7000 series? With that big unknown we won't put too much emphasis on price right now. In a month or two, or six, the graphics card landscape, and what you get for your dollar could look very different. The MSI RTX 4080 Suprim X will remain an excellent card, regardless of where its price ends up.

MSI's Suprim X is designed and built as well as any card we've ever seen, but like every other RTX 4080, it's far too expensive right now.

Chris Szewczyk

★★★★☆

	Metro: Exodus Enhanced Edition				F1 2021
	2560x1440 Extreme Preset	2560x1440 Extreme Preset	3840x2160 Extreme Preset		2560x1440 Ultra High Preset
	Min	Avg	Min	Avg	Min
Nvidia GeForce RTX 4090	76	121	53	77	204
Asus TUF Gaming GeForce RTX 4080 OC Edition	60	94	40	53	173
MSI GeForce RTX 4080 Suprim X	59	95	40	54	175
Nvidia GeForce RTX 3080 Ti	48	67	29	39	131

Asus TUF Gaming GeForce RTX 4080 OC Edition

A big and TUF RTX 4080 – for a price.

■ \$2,599 | www.asus.com

Any discussion involving high end GPUs typically includes Asus ROG Strix and TUF models. The RTX 4080 delivers on performance and power efficiency, but it's still a 320W card, and that means you need a beefy cooler to keep it operating at its highest boost clocks. The Asus TUF Gaming RTX 4080 OC Edition comes with just such a cooler, though the competition is fierce at its price point.

The AD103 GPU at the heart of the TUF is an amazing chip. It's a 4nm powerhouse with a 60 percent higher transistor count over the previous generation GA102 flagship. It brings more cache, more shader cores and improved ray tracing and AI capabilities. It really is an excellent GPU, though for now, Nvidia has priced it too high. Let's hope that changes.

Asus' TUF brand has evolved from what was originally a military and reliability themed range to Asus' mainstream gaming brand, positioned below ROG products. While the 4080 TUF doesn't have all of the distinguishing features of the ROG Strix 4080, it doesn't compromise on cooling. The 4080 TUF is a big card, effectively taking up four slots. If the Strix is the attention grabber, the TUF is the more utilitarian option, though there's a small TUF logo with RGB lighting on the right side of the card.

Despite its large size, it wasn't quite able to match the cooling of the MSI Suprim X, though a recorded peak temperature of 61 degrees while looping 3DMark Speedway is still an excellent result. The card comes with a rated boost clock of 2,595MHz, though it has an OC mode at 2,625MHz. As is almost always the case if your cooling allows, the card was happy to boost beyond that, resulting in 2,730MHz while running 3DMark Speedway.



■ Asus TUF Gaming GeForce RTX 4080 OC Edition; 9,728 CUDA cores; 2,595MHz boost clock; 16GB GDDR6X 22.4Gbps memory, 716.8GB/s memory bandwidth; 3x DisplayPort 1.4a, 1x HDMI 2.1; 320W TDP, 1x 16-Pin power connector.

The TUF includes a welcome extra HDMI output, giving you two HDMI 2.1 ports and three DisplayPort 1.4a ports. That gives you a bit of extra display flexibility.

The RTX 4080 delivers awesome performance, as does the Asus TUF. While the Asus clocked itself a tiny bit lower than the MSI Suprim, the real-world impact is zero. Most games that struggled to break 60FPS at 4K can now do so, even with ray tracing effects enabled, and that's before enabling DLSS. Even in the most demanding titles, 1440p at high refresh rates at the highest settings is possible, all while remaining effectively silent. Even our quiet AIO pump whirr was more noticeable.

The Asus RTX 4080 TUF is an excellent graphics card, but suffers from the same inflated pricing that afflicts all

RTX 4080's, but more so in this case. It costs the same as the MSI Suprim X, which is MSI's flagship card which gets the nod for its premium finish. It's more of a Strix competitor. But if the price of the TUF comes down to a more palatable level, then the whole equation changes. We'll know more in a few months once the market settles.

You know that you'll get a cool, quiet and well-built card when you buy a TUF model from Asus, but while it remains at this high price, you should also consider competing options. The cheaper non-OC RTX 4080 TUF is one such option.

The Asus TUF is a well built, cool and quiet 4080, but it needs a price readjustment when compared to competing flagship offerings.

Chris Szewczyk

★★★★☆

	F1 2021			Ghost Recon: Breakpoint	Ghost Recon: Breakpoint	Cyberpunk 2077				3DMark
	2560x1440 Ultra High Preset	3840x2160 Ultra High Preset		2560x1440 Ultimate Preset	3840x2160 Ultimate Preset	2560x1440 RT Medium Preset		3840x2160 RT Medium Preset		Time Spy Extreme overall
	Avg	Min	Avg	Avg	Avg	Min	Avg	Min	Avg	
	243	138	160	168	119	83	99	42	50	16,810
	201	99	117	145	96	64	76	30	36	13,050
	202	99	117	146	96	62	76	31	37	13,193
	153	76	89	119	81	50	60	24	30	10,142

OLED TV / MONITOR

LG UltraGear 48-inch UHD 4K OLED 48GQ900-B

Is LG's first large format OLED gaming monitor the future of PC gaming or something that shouldn't have left the living room?

■ \$1,834 | lg.com/au



apc
HIGHLY
RECOMMENDED

LG is pitching its UltraGear 48-inch UHD OLED TV as the company's first OLED gaming monitor ... but instead of aiming for the super enthusiast as Samsung has with the Odyssey range, LG is hoping to make a more approachable product. The 48-inch UltraGear UHD 4K OLED Gaming Monitor was on sale for \$1,834 at the time of writing, which is a pretty fair price for a sub-55-inch OLED TV these days. Sure you'll be able to find cheaper OLED TVs at a similar size out there, but they'll come with concessions like a slower framerate, higher response time, no DisplayPort input and a lack of gaming monitor screen optimisations.

If we start with the basics, the 48-inch UltraGear 48GQ900-B is a great monitor. It's expansive, offering a TV-sized gaming experience, but it also comes with the rich colours offered by OLED and 99-percent coverage of the DCI-P3 gamut. These features, plus the 120Hz refresh rate that you can play in HDR modes add up to an amazing looking gaming experience. There's no denying that the level of picture quality is impressive and, when fully engaged, will challenge even the most expensive gaming PCs.

If you do want a little extra frame

■ 47.53-inch OLED display at 3840 x 2160 pixel resolution; 330cd/m2 Typical brightness; Typical contrast ratio 1,500,000:1; 120Hz refresh rate (138Hz overclocked); 99% DCI-P3 coverage; HDR10; AMD FreeSync Premium; G-Sync Compatible; 107x66x18.5cm; 16.8kg.

rate performance you can overclock the screen to 138Hz, which is perfectly serviceable for online competitive titles, but anything over 120fps does mean you'll have to give up AMD and Nvidia's adaptive sync technologies in order to push it that last little bit. That said, it's a sacrifice you might be willing to make for the occasional fast-paced FPS.

\$1,800 might only be enough to get you a mid range TV, but that budget is very much at the top-end for a gaming monitor, so consumers expect a lot more than just picture quality. Unfortunately LG hasn't quite nailed these additional aspects of the 48GQ900-B.

The monitor interface hasn't been optimised for the UltraGear Remote Control scroll wheel so it takes as long as buttons or a joystick to navigate the vertical menus and the controller is unnecessarily large, occupying more space than a large gaming mouse. There's also a distinct lack of options to customise picture quality which made us feel like the display had

unrealised potential.

The monitor stand is really more of a basic TV stand since it doesn't offer any tilt or height adjustment, but it's been designed to look like a monitor arm, which makes it less stable towards the back corners. An intentionally non-VESA 300x200mm mounting point will also limit your options when trying to find a reputable monitor arm alternative.

The other big downside is that it has been stripped of all smart TV functionality, so don't think you'll get any of the streaming apps or digital TV functionality you see on the Samsung Odyssey Ark. With a typical brightness of 330nits it's around double that of an Aorus FO48U Gaming Monitor, but it's still a way off the best TVs you can buy.

An amazingly immersive living room PC display that needs better controls and software.

Joel Burgess

★★★★☆

MONITOR

BenQ Zowie XL2566K

How much do you really like gaming? BenQ's latest Zowie will extract the truth from you.

■ \$1,099 | <https://www.benq.com>



Early last year we reviewed a 24.5-inch Full HD gaming monitor with an old-school TN-technology panel replete with its poor viewing angles, mediocre colour reproduction, modest brightness, mediocre contrast and eye-watering \$899 price tag. It was the 240Hz BenQ Zowie XL2546K and we gave it top marks and declared it revolutionary. Now, here's the next evolution – the similarly sized, \$1,099 360Hz Zowie XL2566K. Well, alrighty...

Firstly, it's worth reminding newcomers that you're not expected to see all 360 frames each second on a 360Hz refresh rate display. What you will notice, however, is just how smooth and accurate all objects in motion look and this is particularly important for competitive gamers who need to see any and all, tiny, rapidly-moving objects as quickly and clearly as possible. If they don't, then the literal split-second it takes for an enemy to appear around a corner might be missed and it's game over. That's a huge deal when major tournaments are on the line.

However, refresh rate isn't enough – you also need fast-transitioning pixels which get reported in spurious grey-to-grey (GtG) transition speeds, the delineations of which are different

■ 24.5-inch, matte, 1,080 x 1,920, flat-screen, 360Hz LCD with TN panel. 320 nit brightness, 1000:1 contrast. DyAc+ anti ghosting technology. DisplayPort 1.4, 2 x HDMI 2.0, 3.5mm audio jack. Tilt -50/+230, Swivel -450/+450, 15.5cm height adjustment. Wired remote. 521-366 x 571 x 200 mm, 6.2kg.

depending on who you ask. No panel's pixels can change fast enough to keep up with refresh rate unless you add extra voltage (called overdrive). If you add none you get ghosting edges trailing objects. Add too much and you get inverse ghosting and overshoot where the edges of objects keep going after motion has stopped. Either way, you get a degree of image smearing and ghosting which ranges from massive to extremely minor (on the fastest panels).

Why the lecture? It's because this combination of pixel response, overdrive and refresh rate is hard to get right on any gaming monitor, let alone at 360Hz. And the news that pro-gamers want to hear is that the Zowie XL2566K has the best combination on the market. Even when viewing in super slow-motion, it's noticeably better than 240Hz monitors.

That's not all. The Zowie uses BenQ's DyAc+ backlight strobing technology which inserts a black frame between each transition to further promote smoothness and clarity. In the past, ramping this up has detrimentally affected brightness, but there's very little

impact here. It results in even more clarity and additional smoothness.

Best of all is that it's easy to tinker with all the settings to get everything just so. In addition to three overdrive pre-sets you can incrementally adjust it with a 0-30 slider plus, unlike rival models, you can adjust brightness while DyAc+ strobing occurs.

Other features include blinker flaps on the side of the screen; a solid, small-footprint stand for getting right up close; generous height and pivot adjustments with notched measurements; carry handle and PVC cover; plus a wired remote.

If gaming is life and you love the immersion of a bright, colourful, HDR display, this isn't it. If gaming is your job, or something you take incredibly seriously, this a professional tool that you need to buy. If the latter sounds like you, then the immense \$1,099 price tag is irrelevant.

In a space where many claims are boastful,
BenQ's new 360Hz panel truly delivers.

Nick Ross


★★★★★

PC CASE

Nzxt H7 Flow

Is this the pick of the bunch?

■ \$219 | nzxt.com



Nzxt's latest H7 line-up of cases consists of three new models, the H7, the H7 Flow, and the H7 Elite. The Elite was an impressive and handsome chassis and, as the most expensive option, we'd naturally assumed that it would be the best performing choice out of the three. But while the main focus of the Elite model was its good looks, how does its sibling, the Flow, compare? Costing \$110 less than the Elite, is it a better option?

Visually, the Flow sits in the middle of the three H7 cases. It's better looking than the standard case with its flat front panel but not as striking as its more expensive sibling, the Elite. However, looks aren't everything. The front panel on our review sample features a matte white perforated cover that allows an unrestricted flow of air through the case. Nzxt has executed this well as it not only looks stylish but also allows light from RGB fans to shine through. The clue may be in the name, but this functionality is the H7 Flow's main selling point.

Accessing the internals of the case is a breeze. Just like the H7 Elite, the Flow has thumb screws to hold the

panels in place. The panels pop off and slide in easily, which is a bonus for those who are planning on making any upgrades or replacing parts in the future. Inside the case, there is more than enough room for either a Mini-ITX, Micro-ITX, ATX, or even an EATX build, so long as the latter motherboard is up to 272mm. This allows for a lot of versatility, making it a great size for most custom PCs.

As well as great compatibility, the H7 Flow features a wide choice of routing options for cable management. Of course, you can pile the cables in, shut the back panel and cross your fingers, but you'll regret it down the line. It's always best to set aside some time to plan your cabling and Nzxt does a great job of helping the process with plenty of channels for your wires. There are also Velcro straps to keep everything neat and tidy.

Unlike the H7 Elite, the H7 Flow comes with only two 120mm quiet Airflow fans. Whereas the former comes with three 140mm RGB fans and one 140mm quiet Airflow fan.

Going back to the H7 Flow, the build quality remains high, just as

with all the H7 models. Without being too heavy, the case feels rigid and has even panel gaps around it, showing that quality control is important to Nzxt. Internally, it features the same clearance and support offered by the standard H7 and the Elite. That's availability for three 120mm fans or three 140mm fans at the front, three 120mm fans or two 140mm fans at the top, and one 120mm or 140mm fan at the back. To help your cooling requirements even further, it's compatible with a 360mm radiator at the top with push-pull, a 360mm radiator at the front, and a 140mm radiator at the rear. Overall, the H7 Flow has some great cooling capabilities.

When it comes to the pricing of the H7 range, it's hard to argue against the Flow being the bargain of the bunch. With the H7 Flow, you get premium styling, a high-quality finish, and versatile functionality.

Best thermal performance in the lineup, bold colour combo and solid build quality.

Sam Lewis


★★★★☆

BLUETOOTH HEADPHONES

Focal Bathys

Seriously stunning in almost every way.

■ \$1,199 | focal.com



"Put them on and you feel beautiful, buoyed up by their detail, insight and clarity, but deploy ANC (you get 'silent', 'soft' and 'transparent' options) and you might find you want a little more – which you can't tweak."

Truly audiophile-level Bluetooth cans are a market segment largely ignored until now. French company Focal has long had an strong range of extremely high quality wired headphones, and now it's just owned the scene with this rather incredible listening delight.

Sonically, the Focal Bathys are more than worthy of their enviable heritage. And aside from the quality wireless listen, chuck in the USB-C cable in DAC mode and you've got high-res up to 24-bit/192kHz too. Oh, and in case it needs to be stated, they're utterly stunning to look at.

The Focal Bathys are not the French audio specialist's first foray into beautiful cans – indeed, the company makes some of the best over-ear headphones we've ever had the pleasure of wearing. But this is Focal's first foray into wireless headphones and the Bathys are expensive. We understand the reasons behind the asking fee. Those patented aluminium/magnesium 'M'-dome speaker drivers – made in France using technologies from the finest Focal headphones – don't come cheap; neither does the backlit flame emblem in the centre of each ear cup – but they're more expensive than the AirPods Max and the Bowers & Wilkins PX8.

■ Wireless over-ear headphones, 40mm 'M' dome drivers, 15Hz-22kHz frequency response, Active Noise Cancelling (ANC), Bluetooth 5.1, 3.5mm audio jack, USB-DAC (with up to 24-bit/192kHz audio resolution), fast charging compatible, Siri support, up to 30 hours battery life between recharges.

Put them on and you feel beautiful, buoyed up by their detail, insight and clarity, but deploy ANC (you get 'silent', 'soft' and 'transparent' options) and you might find you want a little more – which you can't tweak.

None of our gripes regarding the Bathys pertain to their sound, irrespective of the noise-cancellation profile you select – and although the cheaper Sennheiser Momentum 4 Wireless boast double the battery life at 60 hours versus 30 hours, the Sony WH-1000XM5 also come bearing 30 hours of stamina.

The Bathys' companion Focal and Naim app (you swipe right for Focal or left for Naim, like some sort of audiophile dating app) is also likeable and has every necessary feature to aid your listening pleasure, including a five-band EQ tab with helpful presets.

Perhaps confusingly, our issues refer to the build – odd since we've just said how beautiful they are, but stay with us. The on-ear physical buttons feel a bit flimsy and prototype-esque, and we do find them hit-and-miss during use. Also, the

headband occasionally clicks during wear which impacts the sound.

Again, for this kind of sonic prowess, we're prepared to forgive most things – for us, sound is paramount. But the big selling point here is wireless listening, so we do need to point out these minor shortfalls in what is still a glorious-sounding set of over-ears. The Focal Bathys are available now for \$1,199. This makes them more expensive than all of the class-leaders in this space, including the Bowers & Wilkins PX8 (\$1,149), Sennheiser's Momentum 4 Wireless (\$549) and the class-leading Sony WH-1000XM5 (\$499). And let's not forget that the aforementioned Sennheiser over-ears boast double the stamina of the Focal Bathys (and the XM5, and the Bowers & Wilkins PX8) the pricing feels bordering on arrogance... But it isn't – because the sound quality here is exceptional.

Focal's debut wireless over-ear headphones are a veritable sonic delight.

Ben Mansill & Becky Scarrott

★★★★★

WIRELESS GAMING HEADSET

SteelSeries Arctis Nova 7 Wireless

Is it a step forward or a step backward for SteelSeries?

■ \$399 | www.steelseries.com

"What really matters is the performance and, on this headset, gaming audio quality is top-tier."



Among the three-strong Nova lineup from SteelSeries, the Nova 7 sits right in the middle. The three different variants include the Nova 7, the PC-optimised version we have on test here, the Nova 7P optimised for PlayStation, and the Nova 7X for Xbox consoles. Costing just short of \$400, it's not exactly at the budget end of the headset market, so is this new addition worth an upgrade or are you better off grabbing an older SteelSeries headset that offers better value for money?

Unboxing the Nova 7, you not only find the headphones but a range of cables and adapters too. There's a multi-platform USB-C wireless dongle, a 5ft USB-C to USB-A charging cable, a 5ft USB-C (female) to USB-A extension cable, and a 3.5mm auxiliary cable. The headset is also wireless via Bluetooth, offering a healthy amount of connectivity choices.

We aren't huge fans of the Arctis Nova 7's design though, especially compared with some of the company's older models such as the Arctis 9 wireless. There's nothing wrong with the design, it looks modern and slick in the matte-black finish of our review sample, but it's not as alluring as previous models. Unfortunately, it falls

short when it comes to build quality. The plastic feels cheap and scratchy and not what you'd expect from a headset costing \$400. The elasticated suspended headband design has been swapped for a band that clips into place. It's still comfortable but doesn't offer as precise adjustment as before. The retractable built-in microphone sits flush with the earcups and the foam padding on the cups is plush.

On the left, there are controls for the mute function, volume control wheel, retractable mic, and the 3.5mm port. On the right is a volume wheel for game/party chat, the power button, the Bluetooth button, and the USB-C port. As we mentioned earlier, the materials used feel cheap and that's especially noticeable on the contact points, but overall, the Nova 7 headset feels sturdy and is comfortable too.

What really matters is the performance and, on this headset, gaming audio quality is top-tier. One thing we should point out is that it's worth installing the native software, in this case, the SteelSeries GG application. This included a firmware update for the headset, which we installed before testing. In the app, we selected the bass boost EQ and this gave

us the best sound. Where the standard preset lacked a little oomph, the bass boost EQ fixed that.

With the software sorted, the Arctis Nova 7 wireless offers a well-rounded and detailed soundscape. Playing games such as Assassin's Creed: Origins and other open-world titles, the Nova aids immersion without sounding muddy. Directional accuracy in competitive FPS game modes was also pinpoint, another strong point of SteelSeries headsets. The same goes for the microphone, pickup was strong and, on recordings, our voice sounded bright and clear.

If anything, it could offer deeper bass as standard but that can be addressed in the app. Where the headset doesn't live up to the price tag is in the overall design and build quality. For around the same price, you can pick up either the Arctis 7+ or the Arctis 9+, both of which offer similar audio profiles and much better build quality.

Clear gaming-focused audio, a 38-hour battery, great connectivity and comfortable.

Sam Lewis

★★★★☆

GAMING MOUSE

Corsair Katar Elite Wireless

Does this Katar provide razor sharp accuracy or is it a blunt instrument?

■ \$TBA | corsair.com



"It's lightweight, has a silky smooth gliding action, and a rechargeable battery with a long life. It also offers three types of connectivity, either through Bluetooth, USB Type-C to Type-A cable, or via the Slipstream USB dongle"

This Elite variant is the latest and most premium offering in Corsair's Katar mouse lineup. The design doesn't stray far from the range's format, so is this new entrant the best Katar mouse of all?

In the box, there's a USB Type-C to USB Type-A cable and a wireless USB dongle hidden in a small compartment beneath the mouse. These pesky receivers can easily get lost, so we always prefer it when the mouse can house the dongle. The Katar Elite follows a similar design to the most recent Katar siblings, the Katar Pro XT and the Katar Pro. The design favours the claw and fingertip gripping techniques and is symmetrical. It's coated in a soft matte finish, which feels comfortable in the hand but can start to look greasy after a little use.

We prefer this smoother choice of material in comparison to the Katar Pro Wireless with its more textured surface. Speaking of which, on either side of the mouse, there's a textured plastic grip where your thumb and little finger sit. This might be more effective if it was rubberised, but it does a good job of helping you keep control when quickly flicking the mouse.

Another feature that helps maintain full control is the PTFE feet and sensor surround. Paired with the super

lightweight design, this helps it to glide effortlessly. There's virtually no drag when swiping this device around your mouse mat, which makes using the Katar Elite a breeze.

The two primary buttons feature Omron switches and Corsair Quickstrike button technology. This means there are zero gaps between the key and the switch, resulting in low travel and almost instantaneous actuation. This combination makes it an ideal gaming mouse, giving you a seamless response time. The scroll wheel is decent but too light for our liking. The feedback when hitting the increments is soft, meaning that you can unintentionally scroll further than you'd like. Thankfully, it's rubberised and has plenty of grip. There are also two side buttons where your thumb sits, set to go forwards and backward in a browser as standard.

On paper, the Katar Elite's specs improve on those of its predecessor with a higher DPI sensitivity of 26,000. Admittedly, that's impressively fast but the option to change it is there. It also contains a Marksman 26K sensor that, during our testing, performs responsively and accurately. These factors ensure this mouse is aimed towards speed-dependent gamers, such

as those who play competitive FPS titles. As it doesn't feature an abundance of macro switches, we wouldn't necessarily recommend this lightweight mouse for MMORPG players, for example.

Anyway, it's fair to say that this mouse ticks a lot of boxes for competitive gamers. It has a pin-point accurate sensor, rapid DPI, and polling rate options available within the iCUE software. It's lightweight, has a silky smooth gliding action, and a rechargeable battery with a long life. It also offers three types of connectivity, either through Bluetooth, USB Type-C to Type-A cable, or via the Slipstream USB dongle – we'd recommend the latter for the most seamless experience.

If all that sounds pretty good, that's because it is, but a few things let it down. The mouse could benefit from rubberised grips and a stronger scroll wheel, and it also feels cheap due to its lightweight nature. If you can look past these issues, the Corsair Katar Elite Wireless is a dependable gaming and general purpose mouse.

Feather-like weight and 110+ hours battery life, plus pinpoint accuracy.

Sam Lewis

★★★★☆

SMARTPHONE

Google Pixel 7

Properly premium at an appealing price.

■ 128GB, \$999 | store.google.com



The Pixel 7 brings much of the same goodness as the Pixel 7 Pro (APC 514, score 5/5) – in a cheaper package. Costing a very reasonable \$999 with 128GB of storage, it has the same Tensor G2 chipset and the same Android 13 software, which together add up to a great experience.

The design isn't identical to the Pro but it's very similar, with a Gorilla Glass Victus casing, an in-display fingerprint reader and a metal camera band across the back. The display is slightly smaller, at 6.3-inch, which makes it more pocketable, and in our view more comfortable to hold in the hand. The Snow and Obsidian colour schemes return, but in place of Hazel you get the choice of a cool Lemongrass hue.

While the OLED screen isn't quite as stunning as the one on the Pixel 7 Pro, it still looks rich and colourful, with a sharp Full HD+ resolution. It hits a maximum brightness of 1,000cd/m2 in normal

■ 8-core Google Tensor G2 processor; 8GB RAM; Mali-G710 graphics; 6.3in 90Hz AMOLED screen, 1,080 x 2,400 resolution; 5G; 128GB/256GB storage; IP68 rating; dual 50MP/12MP rear cameras; 11MP front camera; Wi-Fi 6E; Bluetooth 5.2; NFC; 4,355mAh battery; USB-C connector; Android 13; 73 x 8.7 x 156mm (WDH); 197g; 1yr warranty.

use, ramping up to a peak of 1,400cd/m2 in bright surroundings. The most noticeable compromise is the 90Hz refresh rate: it's far from jerky, but it doesn't have the perfect fluidity of the 120Hz screens found on the Pixel 7 Pro and other flagship phones.

In use

The Android 13 OS is characteristically clean and easy to navigate. Google's Material You "design language" provides smart colour theming, intelligently generated to complement your chosen wallpaper, while some neat interface enhancements make it

easier to get things done. For example, there's a one-handed mode that moves the UI down the screen to make elements more readily reachable, while

superb haptic integration gives a sense of weight and certainty to touchscreen actions.

Everything feels snappy, too. In the usual benchmarks, Google's Tensor G2 silicon appears about as fast as any other Android phone at around this price. But this is the only one with such close integration between the operating system and the hardware, and you can really feel the benefits. Swiping around Android is flawlessly responsive, while face unlock is perfectly quick and reliable.

There are also some Pixel-exclusive software features that make use of the Tensor G2 chip. Its machine-learning capabilities give a boost to all sorts of tasks, from photo processing to speech-to-text. You can ask the phone to live-transcribe audio content, for consumption with the volume down, while the live transcription tool in the recorder app will make students and note-takers very happy. Additional features are promised in future updates, including "Clear Calling" to

"The design isn't identical to the Pro but it's very similar, with a Gorilla Glass Victus casing, an in-display fingerprint reader and a metal camera band across the back."



improve the quality of voice calls, and an integrated VPN. And alongside the Tensor G2, Google's Titan M2 chip takes care of general on-device security, ensuring your passwords and other data can't be extracted.

Top shots

For photography, the Pixel 7 lacks the 48MP telephoto sensor of the Pro, but it has the same 50MP main camera, and a similar 12MP ultra-wide lens (though without autofocus). There's also the same 10.8MP punch-hole selfie snapper, and 4K video recording at up to 60fps with 10-bit HDR.

With the assistance of Google's excellent image-processing software, the Pixel 7 can capture some really excellent photos. Night mode works faster than on last year's models, and thanks to some smart algorithmic work the Face Unblur feature has now been expanded to full Photo Unblur.

A few creative shooting modes are on hand too: the existing Action Pan and Long Exposure modes add dramatic motion blur either to or around a subject, while the new Cinematic Blur mode brings portrait-style bokeh to video. We noticed though that this sometimes struggled

to separate subjects from the background, causing the artificial bokeh to pop and stutter throughout the frame.

Longevity

The Pixel 7's 4,355mAh battery isn't huge, but it should get you through the day: we ended most days at around 15 percent, with typical usage including streaming, gaming and some extended TikTok sessions. The Pixel 7 can intelligently switch to Battery Saver mode if it doesn't think you'll make it to your next regular recharge time, based on your typical routine, while Extreme Battery Saver lets you eke the most out of your remaining charge by disabling most apps and background services, only leaving the ones you've explicitly whitelisted.

Perhaps our biggest disappointment with the Pixel 7's battery performance isn't to do with longevity but recharging speed. At the maximum charging rate of 21W, it took 30 minutes

to get the Pixel 7 to 47 percent, and 100 minutes for a full charge. For comparison, some OnePlus models can go from dead to 100 percent in just 40 minutes. You'll need to provide your own charger too, as there isn't one provided in the box.

You can also charge wirelessly at up to 20W too, and reverse charge earbuds or smartwatches by resting them on the phone.

There are reasons you might choose to pay the extra for the Pixel 7 Pro – chiefly the

smoother screen and more advanced camera hardware. But for anyone on a tighter budget, the Pixel 7 is an ideal alternative, offering a polished experience and all the benefits of Google's custom hardware at an excellent price.

It's hard not to miss the Pixel 7 Pro's high-end features, but this is a great-value buy and remains a premium device

Ben Mansill & Alex Walker-Todd

★★★★☆


"Perhaps our biggest disappointment with the Pixel 7's battery performance isn't to do with longevity but recharging speed."

SSD

Samsung 990 Pro 2TB SSD

A proper contender and 980 Pro successor.

■ \$439 (2TB); \$265 (1TB) | samsung.com



The Samsung 990 Pro is the much-anticipated successor to the popular 980 Pro. It is poised to be the very fastest drive we've tested, outdoing even the SK hynix Platinum P41 and the Western Digital (WD) SN850X. Initially available at 1 and 2TB for \$265 and \$439 respectively, the 990 Pro will also have a 4TB model out next year. Available bare or with an RGB heatsink, the 990 Pro is using Samsung's new Pascal controller and its new V7 TLC NAND flash, a combination that promises to be faster and more efficient.

The 990 Pro has the rest of the Samsung trimmings, including software support with the updated Magician application and the option for encryption. The drive is designed to fit in a wide range of devices, from desktops to laptops to the PlayStation 5 (PS5) console. As Samsung's new top-of-the-line consumer PCIe 4 SSD, the 990 Pro is designed to take a place among the best SSDs in terms of performance and, according to our tests, it delivers.

The Samsung 990 Pro arrives in the 1TB and 2TB capacities at launch. The 4TB model is coming later, in 2023; it will be nice to see a larger capacity option from Samsung. The 990 Pro

■ 8-core Google Tensor G2 processor; 8GB RAM; Mali-G710 graphics; 6.3in 90Hz AMOLED screen, 1,080 x 2,400 resolution; 5G; 128GB/256GB storage; IP68 rating; dual 50MP/12MP rear cameras; 11MP front camera; Wi-Fi 6E; Bluetooth 5.2; NFC; 4,355mAh battery; USB-C connector; Android 13; 73 x 8.7 x 156mm (WDH); 197g; 1yr warranty.

promises up to 7,450/6,900 MBps, sequential read and write, with up to 1.4 million / 1.55 million read and write IOPS. This is more than competitive, exceeding the previous-gen, 980 Pro on all counts but also promising numbers higher than the WD SN850X.

Samsung is offering a variant at each capacity with a heatsink and RGB, taking a page out of the WD SN850X's playbook. Samsung released version 7.2 of its Magician storage software on October 4th, in advance of the 990 Pro's release. This software now allows for RGB control for the heatsinked versions of the drive. It also enables data

migration and offers diagnostics, firmware updates and driver updates.

The Samsung 990 Pro sports a new controller, Pascal, as a successor to the 980 Pro's Elpis. The Pascal is produced in Samsung's 8nm process node and is still ARM-based. Pascal has significantly higher performance specifications. Samsung states that

this is thanks to an optimised NAND data path through "hardware automation technology" with reduced processing time via a "cache algorithm."


Samsung has certainly made advances with the controller but other components are more familiar. The memory, or DRAM, is Samsung's LPDDR4 which has been used on many of its SSD products. It offers a bit of power savings over traditional DDR.

Performance

The 990 Pro delivers the best results we've seen in 3DMark to date, beating the SN850X. Full Power Mode improves performance even more. The 990 Pro also brings astonishing results in PCMark 10 with an even greater lead than it had in 3DMark. Full Power Mode doesn't offer any gains this time, but the 990 Pro simply devastates its 980 Pro predecessor in this test.

Sequential results indicate that the 990 Pro is better than the 980 Pro and generally it fares well. The exception is at QD1, although performance is still relatively good. Random performance is excellent with reads, particularly in Full Power Mode, but write performance is lacklustre in both modes. The former

"The Samsung 990 Pro arrives in the 1TB and 2TB capacities at launch. The 4TB model is coming later, in 2023; it will be nice to see a larger capacity option from Samsung."



"The Samsung 990 Pro is an exceptionally powerful PCIe 4.0 SSD. It did not disappoint, giving us some record-breaking results and performing well where it mattered."

result is not surprising given the relatively low read latency of this flash at 40µs, an area in which Samsung excels. Random read, low queue depth performance tends to be relevant for a drive's "real world" feel, and the 990 Pro is the best we've yet tested here.

Samsung's rated write performance of this flash has it faster than any other on the market, but there are reasons to limit write speeds. The most obvious reason is to provide better performance consistency, that is sustained performance. Lower writes speeds may also induce less wear on the flash.

Other reasons include thermal and power management. The former can improve lifespan and performance consistency by avoiding throttling.

The latter can mean better drive efficiency, particularly in laptops and consoles. Samsung has emphasised improved thermal and power management with this drive, so this makes sense.

Conclusion

The Samsung 990 Pro is an exceptionally powerful PCIe 4.0 SSD. It did not disappoint, giving us some record-breaking results and performing well where it mattered. Sustained write performance seemed a bit relaxed, but the overall consistency of the drive makes it worthwhile. Samsung updated its software in time for the release of this drive and that includes customisable options for power modes. This drive proved to be efficient and thermally-sound, but it is possible to push it for

higher performance if desired.

This drive is a fine successor to the 980 Pro and, perhaps more importantly, a strong rival for the SN850X. Samsung focused on power efficiency and thermal management with the 990 Pro so that its high performance would work well in a laptop or a PS5, in addition to a high-end desktop.

The top criticism we had for the SN850X was its pricing: launch prices were far too high given the other options available on the market, especially given the downward-trending pricing in the NAND and SSD industries. Samsung's MSRP for the 990 Pro are likewise too high. WD quickly reduced its prices well below RRP, making the SN850X a very compelling option. Samsung will likely sell this below RRP and it really has to in order to effectively compete. It is certainly capable of carrying at least a bit of a premium, though, as it appears to be the all-around champion at this point in time.

The 2TB 990 Pro is, in many respects, the fastest SSD we've tested. This high-end PCIe 4.0 drive excels everywhere it matters, is relatively efficient, and has an optional high-power mode.

Shane Dowling & Ben Mansill

★★★★★

				Samsung 990 Pro 2TB	WD Black SN850X 2TB
Crystal Disk Mark	Sequential Read MB/s			7,153	6,982
	Sequential Write MB/s			6,810	6,676
	4K Random Read MB/s	Queue Depth 1		24,506	22,853
	4K Random Write MB/s	Queue Depth 1		74,645	91,200
3DMark Storage Benchmark	Score			4,687	4,466



MOUSE

MICROSOFT MODERN MOBILE MOUSE*Clickin' little.*■ \$59 | www.microsoft.com

This mouse is probably about as small as a controlling device can get before it becomes unusable. It measures 26x60x107mm (HxWxD) and weighs 78g, which makes it small, squat and light enough to slide into your trouser pocket. While this is really handy when you're travelling around and want something small to pack alongside your laptop, you'll probably want something larger to use with a desktop PC.

The other immediate benefit is the price. For \$60 you'd be hard pushed to find another wireless mouse that matches the high build quality on offer here, even if the Modern Mobile Mouse is relatively pared back, with few extra features. It has the three core-function buttons

on top (left, right and a central scroll wheel), and another button on the base that switches it off when you're not using it.

While many wireless mice come with the option of using Bluetooth or a wireless USB dongle, this model is restricted to Bluetooth only.

This wasn't a problem as it connected to devices quickly and performed responsively. The sensor is fast and accurate enough for most Windows work, and it can operate up to 10 metres away from the PC.

It's powered by two AAA batteries, which should last for around 12 months, though it's not rechargeable unless you use your own rechargeable batteries, and charge them separately as and when you need to.

"For \$60 you'd be hard pushed to find another wireless mouse that matches the high build quality on offer here"

Compact, flat and easy to carry around, this is a well-built affordable mouse that makes the ideal controller for your laptop.

APC

★★★★★

■ Bluetooth 4.2, 1800dpi sensor, Three buttons, 10m range, Headphone socket, 26x60x107mm (HxWxD), 78g, One-year warranty.

KEYBOARD & MOUSE

GOOGLE PIXEL WATCH*Looks like it's time.*■ \$549 | www.store.google.com

Smartwatches can tend to look more like a small phone screen you've strapped to your wrist than a timepiece. But Google's designers have done a great job with the Pixel Watch to create something that looks more like a watch than many of its rivals. At 1.2in across, the 450x450-pixel resolution screen is a lot smaller than the 2in screen you get on an Apple Watch, but this also means it's less bulky.

When paired with any Android smartphone, it displays your notifications that pop up automatically or can be summoned with a swipe up on the touch screen. There are apps, maps, directions, and smart-home controls. It also gives you access to Google's brilliant voice-activated Assistant, which is far better than the version of Siri you get on an Apple Watch.

The Pixel Watch doubles

up as a Fitbit – Google bought the fitness-watch manufacturer three years ago, and has fully integrated the technology into this device. It means you need two accounts – both Google and Fitbit – to get the most out of it, but it's worth this initial hassle. Once set up, it's quick and straightforward to monitor your fitness data, such as how many steps you've taken today or how well you slept the previous night.

Google rates the Pixel Watch for 24 hours of battery life. We managed roughly 26 hours when not using the always-on screen, which reduces battery life. However, such a relatively short lifespan can make things difficult to manage when it needs charging, particularly if you're wanting to wear it in bed to track your sleeping habits.

A great smartwatch for Android users, marrying Android, Google Assistant and Fitbit functions, but we'd prefer a longer battery life.

APC

★★★★★



■ 320ppi AMOLED display, 294mAh battery, USB-C magnetic charging, Bluetooth 5.0, Wi-Fi 4, 12.3x41x41mm (HxWxD), 36g without band, One-year warranty.

SMB MODEM ROUTER

Draytek Vigor2865ax

A Wi-Fi performance focused modem router that supports high speed VDSL2 connections.

■ \$700 | www.draytek.com.au



"The 2865ax model includes ultra fast AX3000 spec Wi-Fi 6."



The Vigor2865 series includes a range of small to medium business routers, with VDSL2 model, WAN port and a top notch Wi-Fi. We are testing the higher end, Wi-Fi 6 equipped Vigor2865ax, but there are also models without wireless, others with VoIP options, or even built-in LTE connectivity. The 2865 series is a multi-WAN router, so a key ingredient is being able to multiple internet connections for improved performance and redundancy.

All models include a VDSL2 modem, which supports 35b Supervectoring, as well as being backwards compatible with ADSL2+. With a suitable VDSL2 connection, the 2865 series can handle speeds up to 300 Mbps. Of course you can also use a non VDSL option via the Gigabit WAN port, or even configure Wi-Fi WAN connections. You also get 5 Gigabit LAN ports, and further connectivity is available via the USB ports, which can support 4G modems.

The 2865ax model includes ultra fast AX3000 spec Wi-Fi 6. This is provided by 2x2 MU-MIMO radios and 2402Mbps of bandwidth on the 5GHz band, while the 2.4GHz band can do up to 574 Mbps. While not the very fastest Wi-Fi available, this still offers excellent performance. For those not needing

■ WAN Port: 1x LAN/WAN, 1x RJ-11, LAN Ports: 5x RJ-45 1GbE, USB: 2x USB 2.0, Wi-Fi: 802.11ax, VPN Tunnels: 32, Warranty: 2 years, Dimensions: 241mm x 166mm x 46mm.

quite so much wireless bandwidth, or wanting to use the 2865 series with existing Wi-Fi APs, there is an AC1300 spec model, as well as one with no Wi-Fi at all.

Multiple high speed WAN connections are not much use without enough performance to actually share it out over a wider network, so importantly the 2865 includes hardware acceleration options. It can be configured in different ways, but makes for a big improvement in overall throughput - especially when running multi WANs. For companies that use multiple SSL-VPN tunnels for remote access, hardware acceleration can give a huge 2.8x boost to throughput.

The Vigor2865 series has more features than we have space, so make sure to check out the DrayTek website for the full list. But some key points are up to 32 VPN tunnels, 16 VLANs, ability to handle 1022 IP addresses, 8 LAN subnets, a robust firewall, filtering and QoS options, failover and load balancing management and Wi-Fi hotspot options. The 2865 also supports managing up to 20 Vigor Access Points, and 8 Vigor routers. Management can be done via

the 2865 interface itself, or using the Draytek VigorACS network management software. With the latest VigorACS 3, the 2865 even supports SD WAN.

Getting the 2865 set up and managing our network of other Draytek devices was fast and easy. Performance of the wired side was top notch, and while exact Wi-Fi speed will depend on the connecting device we managed 549 Mbps at close range, which dropped off to a respectable 399 Mbps at ten meter range. Importantly the 2865ax handled multiple high bandwidth devices at the same time. It will come as no surprise that such a feature packed modem router has a rather dense interface, but it is generally well laid out and quite intuitive to use. Draytek hosts virtual versions of their network gear online, so you can try it yourself at www.draytek.com/products/live-demo.

The Draytek Vigor2865ax brings high performance AX3000 spec Wi-Fi to a router that already offers great features and flexibility for businesses that use a VDSL2 connection.

Lindsay Handmer

★★★★★

Windows 7, 8.1, 10 or 11, macOS or Linux | Free | www.virtualbox.org

VIRTUALBOX 7.0

Test drive an OS.

VirtualBox provides a free way to try different operating systems within your current OS by running them in a virtual PC (VM). It's the most popular tool of its kind, and this new version makes it even more useful.

The biggest change in VirtualBox 7.0 is the addition of support for Windows 11. Previously, when you tried to install Windows 11 in a virtual machine, hardware-compatibility problems prevented you from doing so, even if your PC met Microsoft's system requirements. Although you could beat this restriction by tweaking the Windows registry, it was a major hassle.

Now, running Windows 11 on a compatible PC is as simple as clicking New in the Machine menu, giving the virtual machine a name and browsing

to the ISO file for the operating system. Tick the box to Skip Unattended Installation, specify how much memory and processing power to allocate to the VM, and select the option to create a virtual hard drive. Click 'Finish' and Windows 11 will be added to your list of VMs, where you can double-click it to launch and run it.

Another new feature lets you monitor how much resource virtual machines use. Click the menu button next to a VM and select Activity to view graphs showing its CPU load, RAM usage and data-transfer rate. To see this information


for your host (non-virtual) PC, click the Activity Overview button at the top of the window.

VirtualBox 7 also adds the ability to fully encrypt your virtual machines to protect any data shared in them. However, this can currently only be implemented via the command line so is really an option for advanced users.

Other changes include a new

notification centre for reporting errors, 3D support based on DirectX 11 and an improved Help section that lets you search VirtualBox's user manual.

"The biggest change in VirtualBox 7.0 is the addition of support for Windows 11."




1 To add a new virtual machine to VirtualBox, click the Machine menu and select New or press Ctrl-N. Give the virtual machine a name then find and select an ISO file on your PC.

2 Your virtual machines are listed on the left-hand side of VirtualBox's main window. Double-click one to launch and run the relevant operating system – this will open in a new window.

3 Click the menu button next to a virtual machine to view details about it, browse a log of your activities within it and access the File Manager, which lets you copy files to and from your host PC.

4 Select the new Activity option to display graphs showing the virtual machine's current resource usage. These include CPU Load, which tells you how much processing power it's using.



macOS 10.1 | \$17.95 | coconut-flavour.com

COCONUTBATTERY 3 PLUS

Easily monitor your Apple devices' batteries.

Above: coconutBattery displays much more detailed information about your Mac's battery than macOS.

macOS provides some detail on the health of the battery in your MacBook Air or MacBook Pro, including usage history and an indication of overall health. But it's pretty basic. To find out more about the state of the battery and its performance, you'll need a third-party solution, and that's where coconutBattery 3 comes in.

Available in standard (free) and Plus versions, coconutBattery enables you to keep a close eye on the battery health of your Mac and iOS devices, and log it over time.

As well as the basics like the current temperature, the number of charging cycles and health, coconutBattery enables you to see data like a device's current charge

and full capacity in milliamp hours (mAh). It will also show you data on the health of your Mac's solid-state drive (SSD). For iOS devices that can be connected by USB or Wi-Fi, you can see the average battery temperature over its lifetime, among other data. Data is logged and you can view and print reports of battery health for your Mac and iOS devices.

Once installed, coconutBattery is easy to use. You can choose via its Preferences whether to have it running in the background, and display the remaining battery life as a percentage in the menu bar. If you do, when you click on the percentage figure, more data is revealed. This includes an estimate of the time until the battery is out of charge, its temperature and the number of charging cycles undertaken.

The main interface has three tabs: one for the Mac on which the app is installed, one to view

historical data, and one for iOS devices connected to the Mac.

We particularly like the search field in the Advanced Viewer that enables you to quickly find the parameter you want. The 'Copy to


Clipboard' button is handy, too. We noticed a discrepancy between what macOS told us was the level of charge remaining (83%) and what CoconutBattery said (78%). macOS and iOS set 100% at lower than the maximum capacity of the battery in order to maintain a buffer, whereas

coconutBattery uses raw battery data to calculate it. Both figures are useful only as a guide, the rest of the raw data coconutBattery provides is of more use.

A great, low-cost way to get battery info on your Mac and iOS devices – the Plus version even more so.

Kenny Hemphill

★★★★★



coconutBattery Plus enables you to see detailed battery info on an iPhone connected to the same Wi-Fi network.

Various licences available | <https://ubuntu.com>

UBUNTU 22.10

Michael Reed examines the latest Ubuntu release and finds that it's a solid offering but unlikely to get your heart racing.

The latest Ubuntu (22.10) has made its customary October appearance with the colourful code name of Kinetic Kudu. As it's not an LTS (Long Term Support) release, it will have to be upgraded in nine months, but in our experience, this tends to be a rather pain-free procedure. However, go back a version number and install 22.04 instead if you're looking for a setup that won't need a major upgrade for five years.

As ever, the official, stated target of Ubuntu emphasises the corporate desktop, but the actual user base is broader and includes hobbyist and home users as well. It's a stable desktop to set up for family members who aren't computer experts, just to give one example of where it can be used. It's not massively resource-intensive, but neither is it a particularly lightweight desktop.

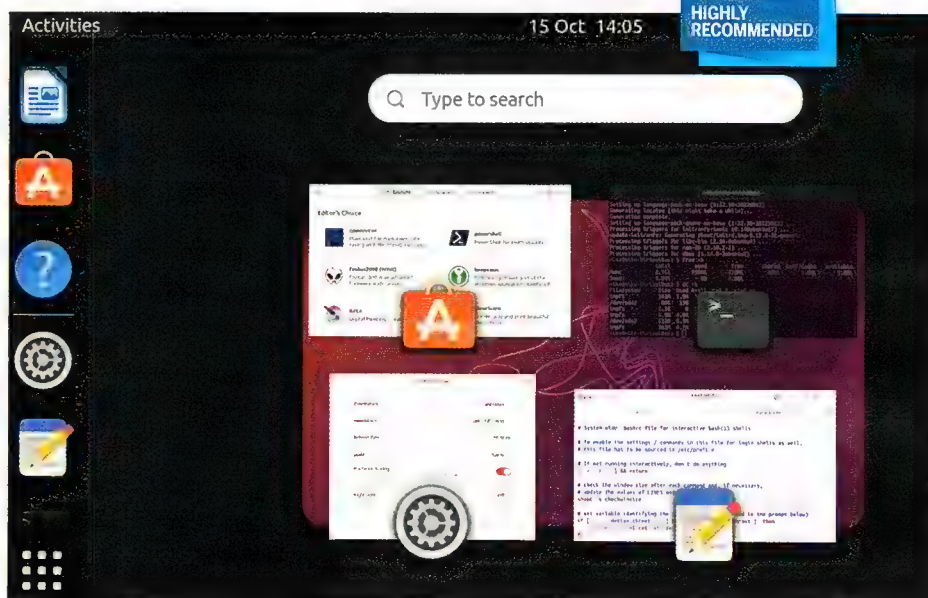
Making it tick

As we have come to expect, this latest release uses a combination of a reasonably up-to-date Gnome desktop on a Debian base. The installation routine is about as easy as it could get and hasn't changed much for several years, and by default it creates a single ext4 partition. Btrfs takes up a bit more CPU time and disk space but offers advanced features such as snapshotting, and distributions such as OpenSUSE have made the leap to it for their standard desktops. As with most of Ubuntu, if you have the know-how, you can choose something more advanced than the default arrangement.

Desktop environment

As this is a Gnome-based desktop, the character of the desktop is closely tied to the current state of Gnome, and this applies to both the layout and functionality, as well as the default applications. On the left-hand side of the screen, we have a simple dock, and placing it down the side makes good use of screen area on a typical widescreen monitor.

One of the easiest ways of adding custom icons to the dock is



by favouriting items in the searchable, pop-up dash. While we were using this dash, it was obvious that measures have been taken to speed up its operation. It wasn't slow by any means in earlier versions of Gnome, but anything that has been done to make it snappier is welcome. It does work well for application launching and locating system settings, but it's missing web search functionality without the use of extensions that have to be installed manually.


Some improvements

In the top right-hand corner of the screen is a status area, and the

functionality here has been expanded, as clicking now brings up a quick settings dialog. From here, you can select the (improved) dark mode, change system volume, access the system settings and lock the screen. You can also invoke the screenshotting tools from here. These are quite good and even extend to starting a screencast video recording, meaning that some users won't have to install applications in these areas. In total, these small improvements are welcome because they add simple access to functions that basic desktop users may need from time to time.

The classic layout is still the same. The dock's at the side and the dash pops up with a keypress.

File Manager has had a clean-up with a disappearing sidebar. It's a simple but functional layout.



Sounding good

Ubuntu now features an upgraded sound system, thanks to the inclusion of PipeWire, a system to share audio hardware between different applications that supersedes PulseAudio. Linux musicians and other audio tech-heads have welcomed this new technology because it has advanced features such as low-latency operation and the ability to pipe audio (and even video) from one application to another. Best of all, for average end users, existing Linux software can take advantage of PipeWire because it dovetails with the existing Linux sound technologies.

Those with long memories might remember the whole debacle when mainstream distributions switched over to PulseAudio, causing many audio-related headaches in its earliest incarnation. However, examining user reports from around the net, it looks as though PipeWire already has a good reputation for reliability, even bringing about improved audio performance, for many, in the area of Linux audio. Most users won't even notice that there has been a changing of the guard on the audio front, and more adventurous ones will want to explore what it has to offer.

In another behind-the-scenes upgrade, Ubuntu now uses the 5.19 kernel with all of the accompanying performance and security improvements that it brings. Hardware support is probably the best that a Linux distribution has to offer. Similarly, Ubuntu must be the most well-documented and supported distro there is, thanks to its popularity.

Too much hand-holding?

Adding packages to the system, we hit upon a little problem, as Ubuntu 22.10 comes with a feature called Unattended Upgrades enabled by default. This checks for and installs important security updates without informing the user. The first problem is that it often badly degrades performance while it's in operation and it can occur at any time. The package management system can't be used while it's working, and it's also impossible to shut down or reboot the system. There is a case for a feature like this on office

As simple as this desktop is, it's extremely easy to add applications with Ubuntu's high-quality software centre.

APPLICATIONS

The application loadout is typical of a Gnome desktop. Expect to find the classics such as the Firefox web browser and LibreOffice office suite. We consider Ubuntu's Software Center to be the best such application in the Linux world, meaning that users can easily add extra applications.

Gedit has been replaced with a text editor that's actually called Text Editor, and it's worth highlighting as an example of typical Gnome application philosophy. In line with Gnome traditions, there is very little to the user interface and it has a level of functionality that one would expect to see in a lightweight Android app. In its favour, it supports syntax highlighting and continuous spellcheck.

It's probably supposed to be used for editing configuration files and similar, but it's only moderately well suited to that role. We tried opening a configuration file in the /etc folder, which it could do, but it uses a proportional font, which threw the formatting out. It was then unable to save changes to the file due to permission problems, and there are text editors that can work around that. Perhaps we're being too hard on it, as it's a recent addition and it may improve over time, but it might be a good idea to replace it with something better from the software centre unless your needs are extremely light.

machines that are left on all day long, as it means that vital patches are installed and can't be dismissed by non-technical users. Looking it up online, there are instructions to disable this feature for users who can be trusted to take care of matters like this themselves when prompted.

Getting the job done

Frankly, this release of Ubuntu has failed to get our blood pumping. It lacks innovation and headline features. Beyond its excellent searchable dash, with Gnome we're talking about a desktop that's less sophisticated than many GUIs that were around 25 years ago.

But how critical can we be of a system that offers a reliable and easy-to-use desktop sitting on top of the best technical underpinnings that Linux has to offer? For business users and non-technical home users, we can confidently recommend an


Ubuntu desktop over pretty much anything else, including the offerings from Microsoft and Apple if you are looking for a basic, reliable desktop for common business-type tasks. If you have enough expertise to operate a word processor and a web browser, you can probably add big-name apps such as Spotify and Zoom if you need them.

More expert users might feel constrained, though. That said, you can install Ubuntu as a base and then add extra applications and even change the desktop environment. Ubuntu is a solid, reliable Linux distro, but it doesn't break much new ground in this release.

Unexciting by default but fully expandable for experts. A few improvements and the best choice for business desktops.

David Rutland

★★★★★





FEATURE PART 1

The (R)evolution of Portable Audio

The new Apple A16
Blowing chip into
the future of
computers.

In this new six-part series, Darren Yates looks back over a century of portable audio, from its humble beginnings to hitching a ride to the moon, from spies to planes, vacuum tubes to the iPhone.

Above: Workers at the Radio Corporation factory in Melbourne put Astor radios together in 1934.



Below: The AWA Radiophone was a 'sealed' radio set sold in Australia in 1923 (Image credit: Museums Victoria/412154/CC-BY-4.0).

Australians have long been technology pioneers and this year, we celebrate the official centenary of public radio broadcasting in Australia, when radio station 2SB – better known today as 'ABC Radio Sydney' – notches up a century of continuous broadcasting in November. Over those 100 years, we've seen audio technology evolve from the mantelpiece to the earpiece, from the radio-valve or 'vacuum tube' to the iPhone. Over the next six parts, we'll look back at how the evolution of technology itself enabled the miniaturisation of audio devices, with the arrival of new inventions that mirrored and changed popular culture. We'll go under-cover, fly to the moon and arrive at what many consider the pinnacle of modern portable audio – the smartphone. But this month, we go back to the 1920s, as Australia discovers radio...


This is radio...

Australians got their first taste of full-time radio on 23 November 1923 when Sydney Broadcasters Ltd powered up their amplitude-modulated (AM) radio transmitter at 8pm and began broadcasting to the Sydney region. However, it wasn't the first station to receive a broadcasting license – that honour went to Farmer & Co, owners of the big Farmer's department store in Sydney. Radio 2FC was likely still building its transmitter and studios inside Her Majesty's Theatre when 2SB got the jump and 2FC didn't begin transmitting until some weeks later.

But what were listeners listening on? In those early days, you were either rich enough to afford an imported set, or you were intrepid enough to build your own (much like PC enthusiasts). But either way, radio sets were only portable in

so far as you could physically lift them, with most early powered models taking up a sizeable space on the mantel piece.

However, not only did you need a license to transmit, you also needed a license to own a radio set – yep, you had to pay for the privilege of listening to radio. What's more, your radio set was also locked or 'sealed' – you could only listen to the stations you were licensed to. Plus, on top of all that, you also had to pay a government license. Radio inspectors would roam the streets on the prowl for illicit radios and those listeners who couldn't produce a license were often prosecuted in the courts. Fines were said to be as high as 20-pounds (approximately \$1,500 today) if convicted. With reports that only 1,400 listeners had signed up by the end of 1923, the system was an unmitigated disaster. The following year, the 'sealed' system was scrapped – radio sets could now tune into any station, but the licensing



system largely remained. Stations at the time were grouped as either 'Class A' (license-funded) or 'Class B' (advertising-funded or 'commercial'). By 1929, all class-A stations were taken over by the Australian government and in 1932, collectively became the Australian Broadcasting Commission (ABC).

The first Australian start-ups

Radio was the first great technology boom in Australia and saw the rise of the first-generation of tech start-ups. In Sydney, Airzone began as a manufacturer of 'no power required' crystal radio sets in 1925 and Kriesler kicked off the following year. Meanwhile, Melbourne saw Radio Corporation start building 'Astor' brand radios in 1926, with Breville launching in 1932.

Still, radio imports into the country were 'big business' at the time and this did little to help the local industry get off the ground – the Australian government decided to help things along, introducing import tariffs in 1930. It had the desired effect with local production rapidly

expanded to meet the growing demand. Dutch giant Philips, for example, opened a local radio factory in Sydney in 1933.

Thermionic or electronic valves ('vacuum tubes' in the U.S.) were as important then as computer CPUs are today and the growing demand for radio meant Australia needed to make its own valves. The Amalgamated Wireless Valve (AWV) company began doing so in 1933 and importantly, launched into development of new power-efficient valves by the mid-1930s. At the same time, events around the world were unfolding that would soon make locally-made components invaluable.

Portable or trans-portable?

Today's CPU technology is manufactured in labs so clean even the dust gets in by appointment-only. But tech in the 1930s was largely hand-made – that meant that most radio sets were at least large enough to sit on the mantelpiece or in many cases, were designed to be eye-catching floor-standing furniture pieces with art-deco style, often with a big 10-inch (250mm)-diameter speaker.



Above left: The AWA Fisk Radiolette from 1936 was a popular mantel-piece radio (Image credit: Museums Victoria/412083/CC-BY-4.0).

Above: A Swedish crystal radio set from 1922, it requires no power source, using the radiowave energy to generate sound in sensitive earphones (Image credit: Holger Ellgaard/CC-BY-3.0).

The valves themselves were roughly 100-150mm tall and around 30-40mm wide – with the better models featuring as many as seven of them, 'small' as we know it today just wasn't possible. However, it didn't mean that portable radio sets didn't exist. Most were portable in so far as they could be moved from room to room. Battery-powered radios were common enough but due to their still comparatively-large power requirements, they needed at least two special batteries – a two-volt 'A' battery to provide the heating power that made the valves work, and a 'B' battery of up to 135-volts to drive the circuitry.

While a 'B' battery could last many months of intermittent use, 'A' batteries required replacement more often.

By the late-1930s, many radio sets were now AC-mains powered and simple enough to use – one knob to turn the power on, another for tuning, one for volume and often a fourth as a tone control.

"Radio was the first great technology boom in Australia and saw the rise of the first-generation of tech start-ups."

War takes radio mobile

When Australia joined Britain and France in declaring war on Germany in September 1939, Australia's technology industry was also mobilised and turned radio from entertainment to military use. But the Japanese attack on Pearl Harbour in



Left: The 12AX7 dual-triode valve, developed in 1947 and still made today (Image credit: Jud McCranie/CC-BY-SA 4.0).



Above left: The Motorola SCR-536 was an iconic WWII-era radio-telephone using new miniature valves.

Above right: The Paraset 7 suitcase radio was used by Allied spies to send Morse Code messages during WWII (Image credit: Timitrius/CC-BY-SA-2.0).

December 1941 arguably had a more salient effect, as did the fall of Singapore and the first air-raid attack on Darwin, both less than three months later. The few radio parts that were then available in Australia for radio makers and DIY hobbyists were now swallowed up by the military and buying any radio set was a difficult task. Local Australian maker Kriesler, for example, had now turned the vast bulk of its radio skills to making communications radio sets for the Royal Australian Air Force.

Australia's work on power-efficient valves during the 1930s was an important step, but for communications to become truly 'portable', those large valves had to be shrunk. In 1941, the giant Radio Corporation of America (RCA) launched the first of a new 'miniature' valve type. Just as Intel and AMD label or group



their CPUs by their respective processor sockets, such LGA1700 and AM4, broader valve families were also known by their sockets. The new 'B7G' valves couldn't deliver the power of older and larger 'octal' types, but they made carrying them light-work. Again, the B7G valves were like mobile-class CPUs – not as powerful as desktop CPUs, but more battery-friendly.

One of the most iconic applications of the new miniature valves during the war was inside the Motorola BC-611/SCR-536 field-radiotelephone – look close enough and you'll see one in most American WWII

movies. They used military versions of B7G miniature valves. Motorola (and other makers under license) reportedly manufactured over 130,000 BC-611s between 1941 and the end of the war. They featured in all of the U.S.'s key European campaigns, from North Africa to Normandy.

Nevertheless, some of the most remarkable design work of the war occurred behind enemy lines.

Winnie the War Winner


One of those designs was created by the signalmen of 'Sparrow Force', an Australian independent company of commandos operating on the island of Timor during WWII.

The day after Japanese planes bombed Darwin on 19 February 1942, its land forces invaded the island and its capital, Dili. Vastly outnumbered, the Australians escaped into the mountains, but without a transmitter to reach Australia and help, their future looked bleak. Signaller Max Loveless had been a radio-technician in Hobart before the war and when a Dutch officer arrived in their camp lugging what turned out to be just a beat-up commercial broadcast radio receiver, Loveless thought there were still enough worthwhile parts there to have a crack at making a rudimentary transmitter to raise Australia.

Their first effort failed, but a second more powerful design, built on a kerosene tin, miraculously transmitted a signal that reached Darwin on the evening of 19 April 1942. It was even more remarkable given the few tools they had – a small tomahawk axe, a screwdriver and a pair of pliers. It

even featured parts scrounged during night-time clandestine raids within sight of Japanese forces. After headquarters had confirmed their identity, the company's requests for boots and ammunition – and very likely, a new working radio – were actioned. This unruly pile of tech parts had saved the day and became affectionately known as 'Winnie the War Winner'. What's more, 'Winnie' made it out of

"Their first effort failed, but a second more powerful design, built on a kerosene tin, miraculously transmitted a signal that reached Darwin on the evening of 19 April 1942."



Left: For many years, a broadcast listener's license was required to listen to radio in Australia (Image credit: Moya McFadzean/Museums Victoria/15789/CC-BY-4.0).

Timor, back to Australia and today, is regularly on display at the Australian War Memorial.

It's an incredible story and one that shows Australian ingenuity at its finest, but given the scale of miniaturisation today, you have to wonder whether a similar feat could ever be achieved again.

Spy Radio

While 'Winnie' was built out of scrounged parts, portable radio-transceivers were an important part of espionage and resistance groups during the war. In some respects, valves are just complicated light-bulbs and like light-bulbs, they were largely made from a glass envelope or tube, which didn't take much to shatter if you dropped them. To improve their ruggedness, military equipment typically used metal envelopes instead of glass.

These appeared in what are often called 'spy radio sets'. One such radio was the 'Paraset', so-called as it was dropped by parachute to waiting agents and resistance groups. These radios were housed in standard rectangle suitcases of the era. However, to keep the weight down, the Paraset also only had enough electronics inside to send and receive the 'dits' and 'dahs' of Morse Code. Voice communication wasn't always an affordable luxury. Nevertheless, the need to have communications with agents behind enemy lines brought the development of portable radios to the fore.

Entertainment goes mobile...

With the war finally over by September 1945, radio technology returned to more peaceful pursuits and like many other areas of endeavour, electronics was about to take advantage of war-time developments. Pre-war radio sets weighed as much as a desktop PC, but in late-1946, a new breed of 'miniature' valves began arriving onto the market – one of the first being the 12AX7. Developed by RCA (Radio Corporation of America), it measured little more than 60mm x 20mm in size, but featured two amplifying elements or 'triodes', a bit like having a dual-core processor, able to do two different things at once. These were known by their 'B9A' socket and in combination with the smaller 'B7G' models from the war, radio


sets of the 1950s became actual 'portable' devices, shrinking from overweight toolboxes to breadbox and more fashionable carry-bag sized units.

Twin-unit or 'dual' valves became ever more popular, reducing the size and cost of new designs. But with the invention of the transistor in 1947, the demise of the valve, whilst not immediate, was all but sealed. Nevertheless, valves are still being made today and are highly sought after by guitarists the world over, one of the most valued (and valuable) still in production is the 12AX7.

... but a new revolution is coming

By 1945, radio continued to be the only source of spoken news and programmed home entertainment in Australia (TV was still more than a decade away).

For those creating that content, there were two choices – have your actors or musicians in a studio performing 'live' or play a recording from a phonograph record. Consumer-grade 10-inch (250mm) records of the 1930s/40s had a rotation speed of 78 revolutions per minute (RPM) and barely three minutes' play time each side. Larger 16-inch (400mm) transcription discs running at 33.3RPM gave 30 minutes per disc, but even so, it's said that listeners could still tell whether a band was 'live' or being




Above: Winnie the War Winner, a radio transmitter built in the jungles of Timor during 1942 (Image credit: Australian War Memorial/060348).


played off a record.

Nevertheless, those '78' records had created pop culture icons of musicians such as Louis Armstrong, Bing Crosby and Glenn Miller. But with the end of the war in 1945, no practical system for high-fidelity delayed transmission of audio yet existed. Or did it?

Next time, we rewind the clock to 1941, pause there and begin playback of a then-secret technology development that revolutionised the second-half of the 20th century, as we fast-forward our way from war-time discoveries to making a trip to the moon and seeing pop-culture get 'wired for sound' and all the change that brought. ■

Below: The Zenith Trans-Oceanic portable radio sold 35,000 units in three months to April 1942.





HOW TSMC CAME TO

A DECADE AGO, TSMC WAS GENERATIONS BEHIND INTEL IN CHIP PRODUCTION TECHNOLOGY BUT NOW IT'S THE UNDISPUTED WORLD LEADER. **JEREMY LAIRD** LOOKS AT THE RAPID RISE OF TAIWAN'S CHIP FOUNDRY.



RULE THE WORLD

Just seven or eight years ago, it was multiple generations behind Intel in chip production technology. 15 years ago, it was a bulk producer of cheap commodity chips with no track record in high-performance silicon, and 30 years ago, it barely existed. So, how has TSMC come to rule the world?

Today, TSMC can claim indisputable world leadership in chip manufacturing with at least one generational advantage over Intel, arguably two. It produces roughly half of the global supply of customer-

designed chips and over 90 percent of those manufactured on cutting-edge processes. Its chips go into not just PCs and phones, but cars, military equipment, medical hardware, and everything else we depend on these days.

TSMC is based in Taiwan, an island territory operating as a de facto independent state without being formally recognised as such by the international community. For the past 50 years or so, China has viewed Taiwan as a rogue element and has officially

maintained the unilateral right to repossess the island. As far as China is concerned, it and Taiwan remain a single nation.

Modern life as we know it depends to a large extent on the output of a single company based in a country that sort of doesn't exist, itself situated on an island separated from China by a sliver of water that might just be the most dangerous place on Earth. How, exactly, did we get here and what does the future hold?

If you're a tech fanatic of a certain vintage, viewing everything through the prism of the PC, you'll probably be familiar with TSMC. But you'll also instinctively think of it as something of a budget outfit, operating slightly behind the curve. Sure, TSMC was critical – it enabled the creation of important PC components, such as graphics chips. Neither Nvidia nor ATI (before the latter was acquired by AMD) produced their own chips. They designed them, of course, but paid third parties, and often that was TSMC, to do the manufacturing honours.

So, TSMC mattered. But GPUs historically did not require the absolute best in chip production technology. Decent transistor density was certainly desirable. But GPUs were not designed for and did not demand the last word in clock frequencies, typically running at a third or less of the speed of contemporary CPUs. In short, TSMC

simply didn't compare with Intel when it came to the fine art of turning sand into ultra-high-performance silicon.

For AMD, increasingly that was a problem. AMD had been forced to spin off its own chip production facilities, also known as fabs or foundries, in 2009 in an attempt to keep the company alive. Building new fabs is a multi-billion dollar enterprise and AMD simply didn't have the cash. The idea was that selling off its existing fabs would give AMD a desperately needed cash injection while allowing the newly independent manufacturing unit to acquire external customers and outside investment, both of which would bankroll the next-gen production units needed to keep up with Intel's world-beating fabs.

The logic was sound, GlobalFoundries was soon the third-ranked foundry in the world and AMD got its cash. Unfortunately, the technology half of the grand plan

didn't work out quite so well. By 2010, Intel's 32nm chip production node was full steam ahead with 22nm arriving the following year. Meanwhile, GlobalFoundries was struggling to get its own 32nm tech working at all. As a consequence, AMD's first quad-core 32nm CPU, codenamed Llano, was delayed until late 2011.

AMD realised it needed to at least think about alternatives to GlobalFoundries. So it turned to TSMC, a company not then known for process nodes capable of supporting high-frequency CPUs. The first 'GloFo' chip of that era, a dual-core APU known as Ontario, was on TSMC's 40nm node and went largely under the radar. Over the next few years, GlobalFoundries got its act together with Samsung's help to create a reasonably competitive 14nm node for 2015. Both the original Ryzen 1000 Series and the Ryzen 2000 follow-up were built on that 14nm tech. But it was a temporary stay of execution

MY NM IS BETTER THAN YOUR NM


On paper, it seems as though TSMC is way ahead of Intel. After all, it's been over two years since the first commercially available devices powered by 5nm TSMC chips hit the market. Meanwhile, TSMC is underway with mass production of 3nm chips. And Intel? Its most advanced node for commercial chips is 10nm. Oh dear!

But hold on. Making assumptions based on the nomenclature involved is problematic. Put another way, TSMC's nanometers are not directly comparable to Intel's. How can that be possible? The name attached to a given process node has become detached from the physical reality of the chips involved. It used to relate to the smallest feature inside a chip. Now node naming is more branding and marketing than science.

In the end, what matters is transistor density or the number of transistors any process node can cram into a

given area of silicon wafer. More transistors mean more features and performance, while smaller transistors use less power. It's a win-win as feature sizes shrink. Such are the advances in that regard, current process nodes are measured in a barely comprehensible metric of hundreds of millions of transistors per square millimeter. Yes, completely crazy.

Anyway, when you compare transistor density, it quickly becomes clear that TSMC's advantage over Intel isn't quite as overwhelming as that 10nm versus 5nm or even 3nm comparison implies. Intel's 10nm node delivers just over 100 million transistors per square millimeter, while TSMC's 10nm clocks in at 53 million. Indeed, TSMC's 7nm node is only good for 97 million transistors per millimeter squared. Suddenly, Intel's rebranding of its 10nm node to 'Intel 7' seems pretty reasonable.



Intel's 10nm node delivers comparable transistor density to TSMC's 7nm.

Of course, TSMC has been making chips on its own 5nm for a couple of years. TSMC 5nm boasts a density of 173 million transistors, well ahead of Intel 10nm.

Meanwhile, TSMC 3nm will up that to a hefty 290 million. Even if Intel manages to get its new 7nm node, branded Intel 4, going next year, it will only edge ahead of TSMC 5nm with 180 million transistors per millimeter squared, placing Intel three years behind TSMC for transistor density.

As for Samsung, its 10nm node was

comparable with TSMC for density. But since then it has fallen behind. And what about the new 3nm node that Samsung proudly announced was already in mass production earlier this summer? It's actually slightly less dense than TSMC 5nm.

As for whether Intel can catch up, well, that's the million dollar – or should that be transistor – question. Intel's 5nm tech, which will be marketed as Intel 20A where the 'A' stands for angstroms, isn't due until at least 2024, putting it 18

months behind TSMC at minimum. Intel's 3nm node, the branding for which hasn't been announced, is expected to deliver a staggering 520 million transistors per millimeter squared, slightly ahead of TSMC's 2nm which is forecast to hit 490 million.

If TSMC and Intel both hit their targets, we can expect 2nm TSMC chips in 2025 and Intel closing the gap to just one year with its 3nm process in 2026. And that's a best-case scenario for Intel. It's a long road back, that's for sure.



Inside Fab 14: TSMC currently has the most advanced chip production technology on the planet.

rather than a long-term solution. In 2018, GloFo announced that development of its next-gen 7nm node was on indefinite hold. GlobalFoundries was, to all intents and purposes, bowing out of the bleeding-edge chip production race.

Strategic alliances

AMD now desperately needed to turn that earlier experimentation with TSMC into a full-blown strategic alliance. And so it was TSMC that got the nod for the Ryzen 3000 series based on the Zen 2 architecture. In hindsight, it seems like the rest is history. But even then, as recently as 2018, betting everything on TSMC being capable of delivering on those critical clock speeds seemed like a serious gamble. There was little hard evidence then that TSMC could produce CPUs running at speeds of well over 4GHz plus, the sort of frequencies AMD would need to take on Intel.

As it turned out, TSMC could. And then some. Those first TSMC-produced Ryzen CPUs ran at up to 4.7GHz, thanks to TSMC's latest 7nm node. Any remaining doubts over the Taiwanese foundry's ability to compete at the bleeding edge vanished. Indeed, AMD has just announced the Ryzen 7000 series will hit 5.7GHz, courtesy of TSMC's leading-edge 5nm tech. There can no longer be any doubt that TSMC can

deliver both density and speed.

While all that was happening, Intel was going through perhaps the most troubled period of its entire existence, with the chronic failure of its own fabs to hit process node targets. The net result of which was that Intel's 10nm node was at least five years late, a veritable eon in the chip industry. Even today, it still hasn't fully transitioned to 10nm technology, while TSMC has been cranking out 5nm chips for years and is just about to mass produce 3nm products.

Of course, Intel and TSMC's production nodes aren't directly comparable. As you can see in the boxout opposite, Intel's 10nm or 7nm technology is very different from TSMC nodes of the same name. But by most objective measures, Intel is somewhere between one and two full generations of chip production technology behind. TSMC's advantage is so acute, that AMD's CEO Lisa Su was


recently able to showcase how the latest Zen 4 CPU cores in the new Ryzen 7000 series processors are half the size of Intel's latest Alder Lake cores. For Intel, that has got to hurt.

If that's a potted history of TSMC's rise viewed through the aforementioned PC prism, it's only a small part of the broader picture of how TSMC came to dominate chip production and become arguably the most important company in the world. To get a better idea of just how monumental TSMC's importance has become, let's put some numbers on it all.

Making the world go round

According to market research firm Trendforce, in the first quarter of this year, TSMC generated over US\$17 billion in foundry revenues. That represents fully 53.6 percent of the total global market in manufacturing chips for customers. Think about it. That's over half of the chips that go into everything that relies on integrated circuits. So, not

Taking the contract for Apple iPhone chip production away from Samsung has arguably been TSMC's greatest achievement.



just PCs, laptops, tablets, and phones. But also TVs, cars, domestic appliances, medical equipment, military hardware, telecoms kit, the works. In other words, pretty much everything that makes the modern world go round. And over half of it depends on the output of a single company: TSMC.

Samsung is next in line with a little over US\$5 billion in foundry revenues, with UMC third at US\$2 billion in revenues. Remember GlobalFoundries, the foundry business AMD spun off back in 2009? It comes fourth with just under US\$2 billion in revenues. Of course, those figures relate to all manner of chips, from cutting edge to commodity chips on legacy processes.

GlobalFoundries, as explained, has pretty much given up trying to keep up with Moore's Law and the rest of the industry, preferring to stick with 14nm and above technology. So, if you break market share down into more finely-grained categories based on specific process nodes, TSMC looks ever more dominant as you approach the cutting edge. In 2020, TSMC owned around 90 percent of the sub-10nm market, with Samsung its only competitor. Even in the 32nm to 12nm space, it had 70 percent of the market. You have to go to 45nm and larger before TSMC's market share drops below 50 percent.


It wasn't always this way. Twenty years ago, there were as many as 20 foundries that could plausibly compete in terms of process node technology. Now there is only TSMC and Samsung, with Intel planning to get in on the action as a customer foundry itself and possibly IBM re-entering the fray with a future 2nm node. The reason is the monumental cost of the most advanced chip manufacturing technology. In 2021 alone, TSMC invested over US\$25 billion in new facilities. Its latest 3nm fab currently being constructed in Shanhua, Taiwan, covers 160,000 square metres – about the same size as 22 football fields.

Of course, TSMC will need more than one 3nm fab in the long run. And if it is to maintain its leadership position, it must then create a new generation of production facilities for the jump to 2nm for 2025, or thereabouts. And on, and on. Very likely, TSMC will invest in excess of another US\$100 billion in advancing its production capabilities before the decade is out. These are preposterously big numbers by any metric. And that's why it will be so hard for anyone, including Intel, to unseat TSMC at the top any time soon.

All of which begs the question of how TSMC got into the position to



NODE TO NODE




A timeline of TSMC versus Intel process nodes



invest on that epic scale in the first place. It all started in 1985 when a Chinese-born engineer named Morris Chang was recruited by the Taiwanese government to overhaul its technology strategy. Chang, who had studied at both MIT and Stanford, was a 25-year veteran of none other than Texas Instruments and undoubtedly a leading expert in commercialising semiconductors. In 1987, Chang founded TSMC with funding from the Taiwanese government as a major shareholder, but also Dutch electronics giant Philips and several other private investors.

Key decisions

Early on, Chang made two key strategic decisions that set the template for modern foundries and enabled TSMC's rapid rise. First, he decided to make TSMC a pure customer foundry. It would only make chips for other companies, it would not manufacture its own designs. Second, he priced access to TSMC's fabs



Is TSMC's 3nm tech in trouble?

THE TROUBLE WITH 3NM

If this so far has largely been a tale of glorious success for TSMC, is that a dark cloud we can make out on the horizon?

We speak of TSMC's upcoming 3nm node. It has been widely reported that 3nm has been problematic for TSMC. Blame the pandemic? It could be a factor. Or is it more that TSMC is sticking with existing FinFET or fin field-effect transistor tech while the competition is moving to more advanced GAAFET or gate-all-around field-effect transistors? Either way, there are a few signs that TSMC 3nm isn't quite on track.

For evidence, we refer you to TSMC's own relentless march of new production nodes. In autumn 2018, you could buy an actual consumer product containing a chip produced on TSMC's 7nm node. By the same time in 2020, there were products with 5nm TSMC chips. Well, here we are at the end of 2022, so that means phones and laptops with TSMC 3nm chips are being assembled now,

right? Apparently not. Right at the end of August, TSMC announced that it would begin mass producing chips on its latest 3nm node in September. If that initially seems to fit with autumn 2022 for 3nm chips, hold on. It takes months of painstaking etching and processing before a bare silicon wafer can be sliced into individual chips.

Most likely, the absolute earliest that actual devices with 3nm TSMC chips could be manufactured would be into December, with retail availability in early next year. And that's if everything from here on goes perfectly.

Meanwhile, TSMC's main rival, Samsung, appears to have stolen a march. In July this year, Samsung announced that it was already mass producing chips on its own 3nm node. Significantly, Samsung has opted for cutting-edge, gate-all-around technology for its 3nm node.

Already, the impact of this mooted TSMC 3nm delay is becoming visible.

Apple's recent M2 chip in the revised MacBook Air retained 5nm technology and has proven distinctly power-hungry by Apple's standards.

Apple's new iPhone 14 Pro has the new A16 Bionic chip using TSMC 4nm technology. The standard iPhone 14, meanwhile, uses the A15 5nm chip previously found in the iPhone 13 Pro. Given the recent TSMC announcement regarding its 3nm mass production only being set to begin, it's a clear target missed by TSMC.

Of course, 3nm may only turn out to be a few months behind schedule. We're a long way from the multi-year delays that beset Intel's infamous 10nm node, now rebranded Intel 7. However, from a company that has executed with such millimetric precision over the past decade, even a short delay to 3nm would be conspicuous.

It would also offer a much-needed sense of optimism to TSMC's competition.

INTEL 2.0

Intel's woes with its 10nm process node have already been well documented. The company's response is known as 'IDM 2.0', a triple-pronged strategy to return Intel to market leadership. Part one is a commitment to fix its fabs and push on to 7nm technology and beyond.

Part two involves using third-party fabs to fill in the gaps in the meantime, given that it will take several years to fix those Intel fabs. And part three? To become a major player itself in the foundry business making chips for customers.

If part one seems like an obvious, if difficult and expensive to execute step, two and three are intriguing, even contradictory. As it turns out, Intel is turning to TSMC for that interim measure of farming out some of its chips to a third party for production. But TSMC will also, of course, be Intel's primary competition as it scales up part three of the plan and attempts to win foundry customers. That Intel-TSMC relationship sure is going to be interesting.

Indeed, early reports suggest it is already strained. Intel has revealed that TSMC will make the integrated GPU chiplet or tile, known as a tGPU, for its Meteor Lake series of processors. Meteor Lake is the architecture that will follow Raptor Lake, the latter being due later this year as a replacement for the incumbent Alder Lake architecture.

Anyway, rumours currently abound regarding delays in Intel's launch plans for Meteor Lake, pushing planned production from late 2022 to early 2023 and then even later into 2023, and in turn the impact all that will have on TSMC. Analysts at Trendforce even reported that Intel's delays had forced TSMC to scale back on its plans for

3nm production rollout in 2023.

Intel has since revealed that Meteor Lake's tGPU will actually be built on TSMC 5nm tech and insists that the CPU family is not delayed but, equally, Intel will not publicly commit to an actual firm launch date. So, the real situation is not at all clear.

If that part of the IDM 2.0 plan has got off to a shaky start, what about Intel's nascent foundry business? Intel has announced Qualcomm and Amazon Web Services as early customers.

Qualcomm will be using Intel's upcoming 20A process node, which will use Intel's new RibbonFET transistors, the first major revision to Intel's basic transistor design since 2011. RibbonFET is Intel's version of the GAA or gate-all-around technology that Samsung recently debuted in its new 3nm node and TSMC won't transition to until its 2nm node around 2025.

Some reports suggest the Intel 20A node is expected to begin ramping in 2024. But since

Intel has yet to ship any commercial products on its 7nm node, marketed as 'Intel 4', that would imply Intel 4 will launch in 2023 and Intel 20A in 2024. While that isn't strictly impossible, it seems incredibly ambitious given the five to six years it took Intel to go from 14nm to 10nm production technology.

Before any of that happens, however, Intel's first major customer will actually be Mediatek. Intel has signed up to produce chips for Mediatek's smart edge devices, including routers, switches, and firewalls. But they won't be on Intel 20A or even Intel 4 or 7. Instead, the chips will be produced on Intel 16, a node we've never heard of before which turns out to be a development of Intel's 22nm technology, which debuted way back in 2011.

All of this makes for intriguing reading, but it doesn't get us any closer to knowing whether Intel can really make a go of its new foundry business.



Graphics chips such as AMD's Navi 21 GPU have long been a TSMC specialty.

well below the competition. The idea was to gain market share and recoup revenues over the full life span of a given process node rather than attempting to make fast money upfront.

Chang was playing the long game. And boy, did it work. By 1992, TSMC was already the world's single largest customer chip foundry. By 1996, the company generated over US\$40 billion in sales, listed on the New York stock exchange, and also awarded dividend payments of over US\$100,000 each to its employees. By the following year, TSMC had an overall capacity for one million wafers a year, each fully etched wafer containing hundreds of individual chips. But that wasn't enough for TSMC's ambitions. That year, the company also committed to a 10-year plan to invest US\$14.5 billion in new production facilities. And that increasing investment and scaling of capacity have continued through to the present day.

As of 2021, TSMC reckoned it had the capacity for 13 million wafer starts per year. Do the rough sums on a typical smartphone chip measuring 100mm squared and produced on TSMC's 300mm-wide wafers, and you can fit over 600 individual chip dies into a single wafer. Subtract a few for inevitable dud dies, call it 500 per wafer once yields have matured and you have the capacity for no fewer than 6.5 billion cutting-edge chips in a single year. Incredible.

Of course, some of the dies TSMC makes for customers are much bigger than a smartphone SoC. Navi21, the GPU found in AMD's top-spec Radeon RX 6900 series graphics cards and currently produced on TSMC's 7nm node, measures 520mm squared. But whichever way those wafers are actually sliced in practice, you're looking at billions of chips.

As for production facilities, the current list of TSMC infrastructure makes for staggering reading. The company runs four so-called

Intel CEO Pat Gelsinger: The old Intel hand behind its new IDM 2.0 strategy.





Fab 16 is just one of numerous multi-billion dollar production facilities TSMC has in Taiwan.

'gigafabs' in Taiwan using those premium 300mm wafers, within a further 300mm facility in mainland China. It has a further four 200mm fabs in Taiwan and two more 200mm factories, one in China, the other located in Camas, Washington. As if all that isn't enough, a further 300mm fab is currently being constructed in Taiwan, with yet another underway in Arizona.

If that's an overview of how TSMC's raw capacity has grown, the missing piece of the puzzle is how it managed to not just catch up to Intel in terms of both transistor density and performance, but actually overtake it by some measures. If you compare the two company's timelines starting around 20 years ago, the two had similar levels of chip production prowess in the early days. In 1999, for instance, both Intel and TSMC made chips on their respective 180nm nodes.

Leading by the node

By 2001, both outfits had upgraded to 130nm technology. However, Intel beat TSMC to both 90nm and 65nm processes by a year in 2003 and 2005 respectively and by 2011 was several nodes ahead, producing 22nm chips. TSMC only moved to its own 20nm process in 2014. What's more, Intel's nodes tend to be one

generation ahead in terms of density at any given measure (see boxout opposite). So, it wasn't until TSMC launched its 16nm node in 2015 that it truly had technology on a par with Intel 22nm.

Two things then conspired to completely upend Intel's technological leadership and elevate TSMC to the position of the world's most technologically advanced chip maker. The first was Intel's chronic failure to bring its 10nm node to market. Five long years spanned the gap between Intel's 14nm and 10nm chips, and that's by the most generous measure that counts a tiny number of dual-core laptop chips in 2019 as the first 10nm commercial products. It was actually at least six years between the two nodes if the measure is multi-core

chips produced on a truly mass scale.

The other major event was TSMC's coup in winning production of SoCs for Apple's iPhone away from Samsung in 2014. The Apple A8 chip, as found in the iPhone 6, was produced on TSMC's then-new 20nm node. In 2015, TSMC shared production of the Apple A9 chip with Samsung. But ever since, TSMC has dominated iPhone SoC production as the exclusive manufacturer for Apple. The significance of that cannot be understated.

For starters, it has guaranteed huge volumes for TSMC's latest process nodes. Every year, almost without fail, Apple has taken over 200 million SoCs for iPhones alone. That's not counting chips for iPads and now also for MacBooks and iMacs with the new M1



TSMC manufactures Apple's world-beating M1 processors on its cutting-edge 5nm node.

and M2 chips, also produced by TSMC. That has allowed TSMC to aggressively target new nodes with the guaranteed knowledge that it has a customer ready to pay for hundreds of millions of chips.

Just as important, Apple is arguably the most exacting customer in the industry. It has pressured and pushed TSMC all the way, demanding the very best and being willing to pay for it.

We will never know for sure, but it's likely that the combination of rock-solid financial footing plus external pressure was what motivated TSMC to make at least one critical technical move. In 2019, it began using EUV or extreme ultraviolet lithography for the first time in its 7FF+ node, which was a revision of its 7nm technology of the previous year.

EUV, of course, is widely regarded as critical in moving below 10nm.

Notably, Intel has yet to use EUV technology for any of its commercially available nodes, despite its 10nm process being at least as dense as TSMC 7nm. Many industry observers suspect Intel's failure to adopt EUV technology is a major factor in its recent difficulties in bringing the 10nm node up to speed. Whatever, TSMC's early use of EUV on 7nm certainly translated into a swift and successful jump to 5nm the following year, leaving Intel further behind and elevating TSMC to a position of undeniable technical leadership. Without doubt, TSMC now produced the world's most advanced integrated circuits.

The future for TSMC looks almost as bright as it has for the last few years. It has become so costly to compete in the cutting-edge foundry market, arguably only Samsung, Intel, and

possibly IBM, can challenge TSMC's dominance. Admittedly there are some signs that TSMC's upcoming 3nm technology is proving a little problematic. However, it will take a monumental mess to allow either Intel or Samsung to catch up, let alone overtake TSMC before the end of the decade. TSMC's advantage is that big.

Ultimately, TSMC's biggest threat isn't other chip makers, it's the potential for political or even military turmoil across the Taiwan Strait, the thin slice of water that separates Taiwan from mainland China. And if something goes wrong in that regard, it won't just be TSMC in trouble. It will be every single person living in the developed world, and indeed beyond, that feels the pain. TSMC is that important. That is both a quite remarkable, but also an acutely worrying, fact. ■

WHEN TWO TRIBES GO TO WAR

It's all too easy to flippantly characterise the contest between the big players in chip design and production in militaristic terms. But is TSMC now so important that it could spark not just financial fisticuffs between corporate entities, but a real-world hot war between the USA and China? That is a sobering thought and one that has become rather more ominous in the wake of Russia's recent incursion into Ukraine. After all, there are many parallels between Russia and Ukraine's relationship and that of China and Taiwan.

At their closest point, just 130 kilometres of water separates Taiwan from the mainland. China, of course, has long maintained that Taiwan remains part of its territory under the One China doctrine. For its part, the US has operated a policy of strategic ambiguity regarding Taiwan. On the one hand, the US does not formally challenge the One China

doctrine. On the other, it maintains informal relations with Taiwan including extensive trade and commercial arrangements.

There's no denying how much is at stake. For the US, allowing China to seize control of Taiwan and, in turn, TSMC would mean a huge technological and strategic loss. However, should China take the bold military step of invading Taiwan, it's far from guaranteed that TSMC's foundries would survive intact and fully operational. Far from needing to be bombed to take them out of operation, fabs are hyper-sensitive to supply chain interruptions or any kind of contamination, never mind the loss of critical staff.

And even if China somehow managed to take control of the island while maintaining a short-term operation of TSMC's fabs, the medium to long-term outlook would likely be grim. To take one example, TSMC along with all other current



Just 130 kilometres separates Taiwan from the Chinese mainland.

cutting-edge fabs, relies on EUV lithography machines produced by the Dutch company ASML. One fairly obvious western response to any Chinese aggression in Taiwan would be to cut off the supply of further ASML machines, which would almost certainly set TSMC back a decade or more.

In many ways, Taiwan and TSMC looks like an intractable problem for both China and the US. Any attempt to forcibly seize Taiwan would almost automatically undermine its economic value to China. But

equally, there is arguably little that the US could do to stop China short of all-out war. And in that scenario, TSMC would be the least of everyone's worries.

With all that in mind, it's not terribly surprising to find that both the US and China are taking major steps to improve their domestic chip production capabilities and essentially sidestep the TSMC problem altogether. The US has the Creating Helpful Incentives to Produce Semiconductors for America Act or CHIPS

Act, legislation that was recently signed into law with unusually bipartisan support and sets aside US\$52 billion to boost domestic semiconductor research and production.

Meanwhile, a recent Biden administration review of supply-chain vulnerability concluded that China is pumping something in the order of US\$100 billion into subsidies for its own chip industry. All of which means that the worst-case scenario of a hot war looks unlikely. But economic conflict is pretty much a nailed-on certainty.

THE NO.1 MAGAZINE FOR PLAYSTATION

PLAY

AUSTRALIA

ISSUE 11 CHRISTMAS 2022
AUS \$14.95, NZ \$14.95
GAMESRADIO.COM/PLAYMAG

→ MASSIVE REVIEW

GOD OF WAR RAGNARÖK

A godly gift for PS5?
Read our expert
verdict inside



→ HANDS-ON

PSVR2

Why this headset upgrade
is the best way to
jump into VR

→ WE'VE PLAYED IT!

RESIDENT EVIL 4

Kick horror in the face!
How this revamp is
better than ever

IRON MAN, BLACK PANTHER,
AND MORE! MARVEL'S NEW
GAMING PLANS GO BIG

MOVE OVER MONSTER
HUNTER – EA'S WILD
HEARTS IMPRESSES

OVERWATCH 2:
IS IT WORTHY OF THE
SEQUEL NUMBER?

→ EXCLUSIVE ACCESS

WO LONG FALLEN DYNASTY

Huge demons, bloody weapons, and lightning-fast duels – Team Ninja talks
going beyond Nioh and Ninja Gaiden with this masocore Chinese epic



→ GETTING SCRAPPY
PULLING NO
PUNCHES

Deep Street Fighter 6 test:
we go hands-on for hours



→ UNDERWEAR: RUINED
THE CALLISTO
PROTOCOL UNSETTLES

How the Dead Space creators'
latest pushes horror forward



9 772653 074004

PRINT POST APPROVED: 100007153 ISSN: 1465-0140

PLAY AUSTRALIA IS THE MUST-HAVE MAG FOR PS5 AND PS4 PLAYERS.

EACH ISSUE FEATURES EXCLUSIVE DEVELOPER ACCESS, HARD
HITTING REVIEWS, AND HANDS-ON WITH THE MOST HYPED GAMES
ON THE HORIZON.

ON SALE NOW

THE TRUE STORY BEHIND THE **IBM** PERSONAL COMPUTER

The industry-creating IBM Personal Computer 5150 turned 40 in October. To mark the occasion, [Tim Danton](#) reveals the story of its birth — and destroys one long-running myth in the process.



No one would have believed in the last years of the nineteenth century that this world was being watched keenly and closely by intelligences greater than man's and yet as mortal as his own; that as men busied themselves about their various concerns they were scrutinised and studied, perhaps almost as narrowly as a man with a microscope might scrutinise the transient creatures that swarm and multiply in a drop of water."

So begins 'The War Of The Worlds' by HG Wells, and while some IBM executives of the late 1970s may have taken offence at this comparison – the world being watched was the explosive growth of personal computers, the eyes behind the microscope belonging to IBM – there is a ring of truth to this parallel. The main difference was that the template of the PC, as established by IBM, would outlast all of the "transient creatures" other than Apple.

There was one other crucial difference: those creatures knew they were being watched. Everyone in the nascent microcomputer industry knew it was a matter of time before IBM would make the leap from building mainframes and minicomputers to PCs. The only question was when.

One popular story goes that the IBM Personal Computer was kicked into action in mid-1980 when Atari sent a letter to IBM's then chairman, Frank Cary, suggesting that it could make IBM's PCs. Rather than fling the invitation into the bin, so the stories go, Cary passed it on to Bill Lowe. Now dubbed "The father of the IBM PC", at that time Lowe was IBM's director of entry systems.

Contemporary accounts suggest this is, at best, a blurring of facts. According to Ray Kassar, then CEO of Atari, the potential partnership was instigated by Bill Lowe. "We had two meetings actually, one in my office and another at my apartment in San Francisco with IBM," said Kassar, quoted in the book *Atari Inc: Business is Fun*. But the discussions never got far,

likely due to Atari's proprietary design and the fact its computers could only output 40 columns.

In truth, Lowe didn't need a memo from Atari to tell him that IBM should be building a new computer; it was something he had been convinced of for years. Why had IBM resisted? As Lowe would reflect in 2007, IBM in the late 1970s was in defence mode, "fighting the Justice Department in the US and fighting legal battles overseas" to protect its hardware and software designs and make sure no rival could service its products.

But Lowe wasn't done yet. In the 1996 documentary *Triumph of the Nerds*, he recalled his subsequent conversation with the IBM chairman. "He kind of said well, what should we do, and I said, we think we know what we would like to do if we were going to proceed with our own product. And he said no. At IBM it would take four years and 300 people to do anything, it's just a fact of life, and I said no sir, we can provide you a product in a year. And he abruptly ended the meeting and said, you're on, Lowe, come back in two weeks and tell me what you need."

Birth of the PC

What Lowe needed, it transpired, was a team of 12 young, dedicated engineers who would work flat out for the next year. "We were selected to go work on a top-secret project," said Patty McHugh, a senior associate engineer on the team who designed the motherboard. "Our mission was to get a product into the market in a year using off-the-shelf components."

This was a radical departure for IBM. "The key decisions were to go with an open architecture, non-IBM technology, non-IBM software, non-IBM sales, and non-IBM service," said Lowe, "and we probably spent a full half of the presentation carrying the corporate management committee into this concept because this was a new concept for IBM at the time."

In particular, the only




Above: From left to right, Frank Cary, Bill Lowe and Ray Kassar.

proprietary chip on the motherboard contained the BIOS (basic input/output system). While that gave IBM some protection against copycats, it was hardly Fort Knox. Even by 1980 it was well established that other companies could legitimately reverse-engineer a BIOS – crack that and anyone in the world could build a computer that was 100 percent compatible with any software that ran on the IBM PC. A platform was born.

Famously, the other key decision was to use "non-IBM software". In August 1980, the month in which the IBM board officially signed off on the project, it was unclear who would be providing this software, but Microsoft was already the frontrunner. Jack Sams was the engineer in charge of software development for the IBM prototype and had plenty of experience working in BASIC: he had spent months wrestling with the language in an attempt to get it working on a minicomputer (the IBM System/23 Datamaster), delaying the project by a year in the process, and had no desire to repeat the same mistakes this time around.

By all reports, Sams liked the cut of young Bill Gates's jib. And it's worth emphasising the "young". While Gates was 24 by this time, he still had the physique and face of an adolescent, meaning Sams initially assumed that Gates was the office boy when they met for the first time. But by the end of their second meeting, he was convinced by the young man's brains and professional manner; Microsoft was a company that IBM could do business with.

To Sams's disappointment, though, Microsoft couldn't provide the CP/M operating system that a business computer would surely need. For this, Gates told him, he would need to meet with Gary Kildall of Digital Research. Gates and Kildall had been working together for some time – they even discussed merging their two companies – so Gates had no hesitation in picking up the phone



to Kildall and arranging a meeting on IBM's behalf. Days later, Sams, along with a couple of other IBM executives, flew up to meet Kildall in Pacific Court, California.

Flying close to the truth

Sams's meeting – or non-meeting – with Kildall is the stuff of Silicon Valley legend. Famously, Bill Gates allegedly described it as the day “Gary went flying”, leading to the apocryphal idea that Kildall spent his day piloting his private plane for pleasure rather than meet with IBM.

There are some things we know for sure. Kildall wasn't there at the start of the meeting, with his wife Dorothy greeting the IBM contingent; this makes sense as she handled the company's business dealings. When IBM handed over a non-disclosure agreement that one Digital Research would later describe, euphemistically, as “unidirectional”, she called in the company lawyer.

While Sams initially stated that he never met Gary Kildall, it now seems certain that the Digital Research founder returned from his business trip to Oakland (having flown there; that much appears true) and, after some deliberation, signed the agreement. Things didn't go any

more smoothly from this point on, though, with IBM's stance being that it wanted to buy the rights to CP/M outright for \$250,000. Kildall said no; he wanted to keep to the \$10 per licence agreement he had elsewhere.

Whatever the truth of that day, we know that Sams was not impressed by Digital Research. He wanted to deal with Microsoft, and the businessman in Gates sniffed an opportunity greater than his loyalty and friendship to Kildall. Although “friendship” may be overstating it: Kildall would write in an unpublished manuscript (tinyurl.com/APC515kildall): “Our conversations were friendly, but, for some reason, I have always felt uneasy around Bill. I always kept one hand on my wallet, and the other on my program listings.”

Development of DOS

Those instincts were probably correct. When Gates heard about an operating system called QDOS – the Q and D standing for quick and dirty – he realised he could provide what IBM wanted without paying Kildall a cent for licensing. He also realised that the company that owned the IBM PC's operating system would have immense power; how much better for that power to be in Microsoft's hands than, say, Seattle Computer

Above: Did Gary Kildall really go flying rather than meet with IBM?

Products, creators of QDOS.

It helped that Microsoft and Seattle Computer Products already had a business relationship. Its most important programmer, Tim Paterson, had helped Microsoft develop an add-in card for the Apple II that would run CP/M. It was Paterson who created QDOS for Intel's new 16-bit 8086 processor, which Seattle Computer Products would soon rename 86-DOS.

But there were a couple of problems. First, QDOS was heavily based on CP/M; although Paterson didn't have access to the Digital Research code, there was publicly available documentation that allowed him to effectively mimic its way of working. Second, in August 1980, QDOS remained rough and ready, although it did include some improvements on CP/M. And third, it was owned by Seattle Computer Products: Gates wanted Microsoft to own all the software rights.

Despite these hurdles, Gates and the newly hired Steve Ballmer felt enough confidence in the

WHY THE IBM COMPUTER WAS THE BEGINNING OF THE END

For veteran *Guardian* journalist Jack Schofield – sadly no longer with us – the end was written in plain letters once IBM created the Personal Computer. It highlighted a flaw in the business plan of every single computer maker of the 1980s: to keep going as a company, you needed to have hit after hit after hit.


“Commodore had a hit with the VIC-20, and then a bigger hit with a Commodore 64,” he said in late 2019. “And after that it was bust. So I predicted that all of these companies would go bust because they couldn't hit a winner every single time. They might get one, they might get two, they might even get three in a row. But ultimately, they were all doomed because they couldn't get ten in a row.”

What's more, companies didn't help themselves by their lack of backwards compatibility. “The logical inference at the time was that the IBM PC was going to sweep the world, because you didn't have new blockbuster machines, you just upgrade the old ones. And in fact, out of Hong Kong at the time, there was a continuous stream

of improvements via plug-in cards and add-ons that gradually got incorporated into the build, as it were. And this strategy of continuous improvement meant you could run the same software, because, you know, companies were willing to support software for a successful machine.”

The IBM PC was also a sign of an industry growing up. “When we started out, there was kind of a novelty in owning a computer,” recalled Schofield, “and you were expected to write your own little programs. But once games became established, the computer became just a vehicle for running software. So if you didn't have software, you didn't have a computer that was worth anything. And you didn't have software unless the software houses believed in the future of your machine.”

This was particularly bad news for the British computer manufacturers if they insisted on creating proprietary systems. “All of the British machines seemed to me to be doomed because they had no future development prospects.”



Mark Dean

Senior Associate Engineer, PC Team
IBM Fellow

operating system – and their ability to acquire it – that they could act. They flew out to Boca Raton, home of IBM's new PC division, and pitched for a deal that would make both of them billionaires.

The most important element was ownership. Microsoft would own the DOS to run the IBM personal computer (at that point called Project Acorn, with no one involved aware of the fledgling British company's existence) but license it to IBM for a one-off fee. IBM would not be able to create its own version of the DOS, with any amendments needing to go through Microsoft. Most crucially of all, this wasn't an exclusive deal: Microsoft could also license the operating system to any other companies that came along.

"The key to the structure of our deal was that IBM had no control over our licensing to other people. The lesson of the computer industry in mainframes was that over time people build compatible machines or clones or whatever term you want to use," said Bill Gates in *Triumph of the Nerds*. "And so we were hoping that a lot of other people would come along

and do compatible machines."

As history – and Microsoft's share price – reveals, not only was Bill Gates's prediction right, but IBM said yes. It helped that there was a personal connection: Bill's mother, Mary, had served on the board of directors of charity United Way, as had incoming IBM chairman John Opel. "Oh, is that Mary Gates's boy's company?" Opel asked when he heard of Microsoft's involvement. The deal was signed in November 1980.

There was still much to do. First, Microsoft had to acquire QDOS. With Seattle Computer Products struggling financially, the \$50,000 Microsoft offered was too tempting to resist (needless to say, Paul Allen, who negotiated the deal, mentioned nothing of IBM). QDOS creator Paterson would soon join Microsoft to further develop the code, which needed urgent attention: according to Gates, even on the date he signed the agreement it was apparent that they were three months behind schedule. For the next nine months, he would drive his team of around 40 programmers to the edge with countless all-night sessions to meet IBM's deadlines.

Above: Mark Dean was one of the 12 IBM engineers who made the hardware.

Dirty dozen

Meanwhile, in Boca Raton, IBM's team of 12 engineers were putting in similar shifts under the leadership of Don Estridge, who had taken over the project after Lowe had been promoted. The team's existence, well away from the IBM mothership, allowed them to operate in a very non-IBM way: there was no big corporation clock-on, clock-off mentality here, with engineer Mark Dean describing the team as a "tight-knit family". "We would celebrate together and eat together," he said in the IBM Centennial Film 'They Were There'. "I don't think that any of us slept together but we would do just about everything else together. We would be at work late. We trusted each other."

The hard work paid off, with Patty McHugh stating that the team shipped the first prototype to Microsoft by Thanksgiving (27 November in 1980). With Estridge and Gates in constant correspondence, using an early form of email and meeting in person, the fast pace of development continued.

Nor did the rush mean poor quality. If anything, the IBM computer was the most thought-through microcomputer design yet seen, with an attention to detail and reliability that lesser companies simply couldn't match. For instance, this was the first microcomputer to run a series of hardware tests (POST) during its boot-up sequence; it meant it took longer to start, but better for a slight delay at the start of the day than a hardware error to cause a crash during a crucial calculation.

This was also the first microcomputer to include a parity bit in the memory, which meant the hardware could detect corrupted memory before it caused potentially fatal errors in

Loading CPM.SYS...

CP/M-86 for the IBM PC/XT/AT, Vers. 1.1 (Patched)
Copyright (C) 1983, Digital Research

Hardware Supported :

Diskette Drive(s) : 3
Hard Disk Drive(s) : 1
Parallel Printer(s) : 1
Serial Port(s) : 1
Memory (Kb) : 640

D>a:
A>dir
A: PIP CMD : STAT CMD : SUBMIT CMD : ASM86 CMD
A: GENCMD CMD : DDT86 CMD : TOD CMD : ED CMD
A: HELP CMD : HELP HLP : SYS CMD : ASSIGN CMD
A: FORMAT CMD : CLDIR CMD : WRTLDL CMD : BOOTPCDS SYS
A: BOOTWIN SYS : CPM H86 : WINSTALL SUB : PD CMD
A: WCPM SYS : DISKUTIL CMD
A>

Left: By the time CP/M for the PC arrived, MS-DOS held all the aces.

"The IBM Personal Computer cost a similar amount to the Apple III. At launch, you could buy an IBM PC 5150 with 16KB of RAM, keyboard, and monochrome monitor for \$1,565."

processing. According to Estridge, IBM even paid attention to the contrast between the screen and the monitor's bezel to reduce eye fatigue. The end result was a solid computer that businesses could trust.

Launch day

Was it exciting? To those in the know, yes. The day after IBM announced what it called the IBM Personal Computer 5150 and everyone else called the IBM PC, *The New York Times* quoted Christopher Morgan, editor-in-chief of *Byte* magazine, as saying, "it's one of the most important announcements we've seen in the industry". Michael McConnell, executive vice president of retailer ComputerLand, made this prescient point in the same article: "People will now know that personal computers are not a fad or a flash in the pan."

Apple, which shipped almost 80,000 computers in the US in 1980, was a little more snuffy. In reaction to the IBM Personal Computer, it took out a full-page ad in the *The Wall Street Journal* with the headline, "Welcome, IBM. Seriously."

Tandy, which dominated the US microcomputer market in 1980 with sales of almost 300,000 TRS-80 computers, appeared just as relaxed. "I'm relieved that whatever they were going to do, they finally did it," said the firm's chief of financial planning, Garland Asher, in *The New York Times*. You can almost hear the sneer in his voice as he continued: "I'm certainly relieved at the pricing. They haven't introduced anything that's going to rewrite the ground rules."

The IBM Personal Computer cost a similar amount to the Apple III. At launch, you could buy an IBM PC 5150 with 16KB of RAM, keyboard, and monochrome monitor for \$1,565. That price included BASIC, but you were expected to buy PC-DOS – IBM's licensed version of Microsoft's DOS – for \$40.

After Kildall threatened to sue

IBM due to the similarity of PC-DOS to CP/M, the companies agreed that users could buy the Personal Computer with a version of CP/M (called CP/M-86). This arrived on the market six months after the computer's release, but it was doomed from the start. For one, it cost \$240, six times the price of PC-DOS. More crucially still, by this time there was a flourishing market of software written for PC-DOS. Even at launch, you could buy popular spreadsheet VisiCalc, EasyWriter, and a number of accounting packages. For a bit of fun, you could play Microsoft Adventure too.

Incredibly for IBM, a company famous for taking years to create new products, the first IBM Personal Computers started shipping in October 1981. That's 14 months after it received the green light from the board. It also chose to sell the computers through retail stores such as ComputerLand and Sears. IBM even created a series of ads aimed at consumers – the first time in its 70-year life that IBM had communicated to anyone other than businesses – featuring Charlie Chaplin's 'The Tramp' character.

Modest beginnings

While sales were initially modest compared to the likes of Apple and Tandy – 13,000 units by the end of 1981 – this was due to production limits rather than a lack of demand. *Byte* reported that 40,000 were ordered on the day the 5150 was announced, and IBM spent the next two years desperately trying to catch up with demand. According to market research firm Dataquest, it sold 156,000 computers in 1982. It's easy to argue that this success brought credibility to computers for the first time; a

credibility enhanced yet further when *Time* put "The Computer" on the cover of its traditional "Man of the Year" issue in January 1983.

Tech journalist Dick Pountain was one of the first to buy an IBM

Personal Computer. "I knew it was going to be it," he said. "At this point I'd had an Apple IIe at home, on loan, and didn't like it at all. I didn't like any of the word processors on the Apple and I didn't like the screen. When I saw that MS-DOS was more or less CP/M, and I knew what the IBM name was going to do for the software base of it, I just thought, that's the way to go. And it was."

IBM's honeymoon period would continue for several years to come,

but it's notable that within a year Compaq released a computer that was 100 percent compatible with the IBM Personal Computer. That meant that any software that ran on the IBM would run on Compaq's machine. And, shortly afterwards, those of Dell, Gateway, HP and Packard Bell. As luck would have it, a company called Microsoft, run by a beaming man-child, was more than happy to sell you the software to make it run. ■

Welcome, IBM. Seriously.

Welcome to the most exciting and important marketplace since the computer revolution began 35 years ago. And congratulations on your first personal computer. Putting real computer power in the hands of the individual is already improving the way people work, think, learn, communicate and spend their leisure hours. Computer literacy is fast becoming as fundamental a skill as reading or writing. When we invented the first personal computer system, we estimated that over 140,000,000 people worldwide could justify the purchase of one, if only they understood its benefits. Next year alone, we project that well over 1,000,000 will come to that understanding. Over the next decade, the growth of the personal computer will continue in logarithmic leaps. We look forward to responsible competition in the massive effort to distribute this American technology to the world. And we appreciate the magnitude of your commitment. Because what we are doing is increasing social capital by enhancing individual productivity. Welcome to the task.



Above: Apple cheekily welcomed the arrival of the IBM PC with this ad.

This article is an extract from 'The Computers That Made Britain', written by APC contributor Tim Danton, which was published in May 2021. This 300-page hardback book tells the story of 19 computers that inspired a generation, including the Commodore 64, ZX Spectrum, Amstrad CPC 464, Dragon 32 and Commodore Amiga. You can order it from tinyurl.com/APC515Tim.





Find the best prices
on the latest tech
getprice.com.au

GETPRICE 



UNDER- STANDING

Networking technology needn't be complex. Darien Graham-Smith explores the protocol that underpins almost all networks, from your home LAN to the internet.

IP

When we hear a thing of network such as "you can't be a C", some of us know that we're looking at an IP address. But what exactly is an IP address, and why does it take that form? There are things you might not know – and that's fine, because you don't need that level of knowledge to keep your home or small office network running smoothly and securely.

So even, this information can be useful when it comes to troubleshooting and optimising your network – and it's interesting, too, in a geeky sort of way. Here's a quick guide to the basics of IP and how it works.



1 WHAT NEEDS AN IP ADDRESS?

The full name of IP is “internet protocol”, but it’s not just used for internet connections. In fact, any device that’s on any sort of network must have an IP address. If it doesn’t have an address, it’s not on a network, because nothing else can talk to it.

To be fair, that’s a simplification. There are numerous networking systems that don’t use IP addresses: in the 1990s and before, many firms ran their networks on Novell NetWare, which used the IPX/SPX protocol; others used the NBF system (also known as NetBEUI) popularised by IBM. Modern technologies such as Bluetooth and Zigbee also allow devices to communicate without using IP. Generally speaking, though, IP has become the universal default for networking, because it’s open, mature and

suitable for both local and internet connections.


There also doesn’t have to be a one-to-one correlation between devices and IP addresses. There can’t be two systems with the same address on a network; this would make it impossible to route packets correctly, and may end up making them both inaccessible. However, one machine can have more than one IP address. Virtual machines often have their own addresses, separate from their host systems. And a single computer will automatically acquire two IP addresses if you enable two network interfaces at once, which you might do for traffic filtering or just so you can access two separate networks simultaneously. It’s a fair bet, though, that the vast majority of connected devices in your home have one IP address. The exception is your router – but we’ll get to that in a moment.

2 WHAT DO THE NUMBERS MEAN?

An IP address is normally expressed as a series of four numbers from 0 to 255 (at least it is under the IPv4 standard, which we’re focusing on here). Believe it or not, this convention is intended to make addresses easier for humans to work with: for a computer it’s just a convenient four-byte chunk, or a single 32-bit number.

Since each of the four sections is an eight-bit number, the parts are known as octets. On most home networks, the first three octets will be the same for every device. The last number uniquely identifies each client; self-evidently this means that you can have up to 255 devices on a single network, minus a few addresses that are reserved for special functions.

If you need more than this, you



can create additional subnets. For example, many home networks have addresses in the range 192.168.0.x; if you need more, you can start using 192.168.1.x, and so forth. If you use all available combinations of the third and fourth octets, you get a generous total of 65,025 addresses.

3 WHO DECIDES WHAT GETS WHAT IP ADDRESS?

Every device or server that's directly connected to the internet needs to have its own unique IP address. This will normally be assigned by the ISP providing the internet service, which acquires the rights to address ranges from a regional internet registry. At the top of the tree, the Internet Assigned Numbers Authority (IANA) monitors and manages the registries and makes sure that no-one is using addresses they're not entitled to.

THE DREADED DOUBLE NAT

NAT is a good thing, so double NAT must be better, right? Not really.


Double NAT normally happens when you connect a device with router functions to a network that already has a router – for example, you might find that the only way to get a new Wi-Fi 6 router on to the internet is to plug it into your old ISP-provided router.

A double-NAT configuration means that you end up with two devices handing out IP addresses. This isn't a disaster. The secondary router won't assign addresses that clash: if the old router is issuing addresses in the 192.168.0.x range, the new one might use the 192.168.100.x range. And the two routers should be able to pass packets between themselves so that all systems can talk to one another and get on to the internet.

Double NAT can create problems, though. Since the secondary router doesn't have its own external IP address, nothing connected to it can receive incoming connections. UPnP won't help, as the computer running the application isn't directly connected to the router that's connected to the internet. Similarly, QoS features that aim to prioritise particular clients or traffic types can get confused, since the primary router can't probe devices connected through the secondary one. And while NAT itself doesn't open up any network vulnerabilities, it does mean you need to stay on top of security settings and patches for two routers instead of one.

There are a few possible solutions. Many routers and mesh systems offer a bridge mode, which effectively turns them into network extenders for the primary network: this lets you use their wireless capabilities, but advanced features such as parental controls or network security scanners are normally disabled.

Alternatively, you can manually configure port forwarding to direct incoming traffic on selected ports from the primary router to the secondary one. Again, though, you'll need to configure this twice, so that the secondary router forwards connections on to the appropriate client.



Above: You can easily change the DHCP settings on your router. Above left: These days, IP settings are usually assigned automatically.

However, the individual computers inside a home or office network don't normally have a direct internet connection. They have access to an internet gateway, but their own addresses are internal ones that are only valid inside the local network.

How are these internal addresses determined? When IP-based networking was new, it was common for the network administrator to assign a permanent ("static") IP address to each computer on the network. However, this approach involves a considerable management overhead, and doesn't permit simple "plug and play" connections to the network.

Nowadays we use a system called the Dynamic Host Configuration Protocol – DHCP for short – which allows the router (or a dedicated server) to automatically allocate a unique internal IP address to each device as soon as it connects to the network.

While this makes life easy, it can create confusion, because DHCP allocations don't last forever. This doesn't mean things will get kicked off your network – if a device is still in use when its address lease expires, it will be quietly given a new one, and normally you won't even know it happened. However, with something like a NAS appliance or printer, you might want to ensure that the IP address never changes, to make certain that all connections and network functions will keep working.

There are two ways of achieving this. One is to configure the IP address manually on the device itself, just as you would have on an old static IP network. Just make sure that your chosen address doesn't clash with anything else on the network; an easy way to do this is to pick an address like 192.168.0.3, then go into the DHCP settings on your router and set the DHCP address range to start at 192.168.0.4. Alternatively, you can keep whatever IP address was automatically assigned, and tell your router to permanently reserve that address specifically for that device. You can normally do this from the client list in your router's interface, or using a dedicated IP address reservation page. Strictly speaking, the router is associating the IP address with the hardcoded MAC address of the client's network

interface, so if you were to take the network card out of one machine and put it into another, the IP address would go with it.

4 WHAT'S NAT?

We mentioned above that all the devices on your network use internal addresses that are only valid on the LAN. These addresses aren't picked at random: there are a few address ranges that are reserved for this role, with the 192.168.0.x range being one of the most popular. Thus in my home the address 192.168.0.17 refers to an iPad, but in yours the same address might be used by a smart TV or an Amazon Echo Dot. Since these local addresses are only used inside our respective LANs, there's no possibility of a clash.

The router also has its own internal address (typically 192.168.0.1) – but as part of your home internet service, your ISP also provides your router with a globally unique address. This allows the router to send and receive data to and from other internet hosts, such as websites and online services.

It also allows the router to provide internet services to all the other devices on your LAN; when a computer on your network wants to connect to (say) wikipedia.org, the router forwards the request over the internet. The remote server sees the connection as coming from your router's external IP address, but this isn't a problem: it sends data back, and the router in turn passes it back to the originating system. The latter process, whereby the router takes care of directing internet packets to and from the devices on your LAN, is called Network Address Translation, or NAT for short.

5 SO HOW DO I ACCESS MY DEVICES OVER THE INTERNET?

NAT works brilliantly when the connection is initiated by a computer inside the LAN, but it doesn't work if an external computer is trying to connect over the internet to one of your devices, because in this scenario your router has no way of knowing which internal IP address the traffic is intended for. This makes it impossible for attackers to hack into your computers and home devices by trying IP addresses at random. However, it's a problem if you want one of your systems to be

WHATEVER HAPPENED TO IPV6?

The implementation of IP in use today is actually the fourth version of the standard, commonly known as IPv4. It was introduced in 1982, and at that time the four-octet address format – representing a total of 4,228,250,625 unique addresses – doubtless seemed quite generous.

As internet usage exploded in the 1990s, however, it became clear that sooner or later IPv4 was going to reach its limit. The architects of the internet foresaw that at some point every possible IPv4 address would be allocated, and it would become impossible to connect any more devices to the internet.

Clearly a new addressing system was needed. Version 5 of IP had already been designated as a system specifically for video streaming (although it never caught on), so in 1998 a new general-purpose system was drafted under the name IPv6.

IPv6 is similar to IPv4, but instead of four octets it uses a series of eight 16-bit numbers (sometimes called hexets), conventionally written in hexadecimal form with the numbers separated by colons, so a valid address might be 0201:0db8:0000:0006:00ba:1234:5678:0090. To make things slightly more readable, leading zeros can be omitted and zero values can be skipped altogether, so that address could also be written as 201:db8::6:ba:1234:5678:90.

This arrangement might not be snappy, but it allows IPv6 to support 340,282,366,920,938,463,463,374,607,431,768,211,456 different addresses – enough to allow each person on earth to own 44 octillion connected devices. It's safe to say that everyone who wants an IPv6 internet address will be able to get one, and it even makes NAT obsolete: on an IPv6 LAN, every device in your home or office can have its own globally unique internet address, with the router simply acting as a gatekeeper.

If you're wondering why we don't already have IPv6, the answer is that we do – it's been built into Windows since Windows XP SP1. The trouble is that IPv6 isn't directly interoperable with IPv4, and the cost and upheaval of migrating online services, home appliances and so forth has so far ensured that IPv6 remains a niche technology. Various measures have been taken to keep IPv4 up and running for as long as possible. The internet authorities have been reclaiming and reissuing IPv4 address blocks that were previously issued to organisations that are no longer in existence, while some companies have been buying up unused addresses from businesses that no longer need them. In the long run, though, internet providers will inevitably have to start offering IPv6-only connections, and the dominoes will necessarily fall from there. And while this may mean short-term pain, it will also bring benefits. IPv6 is more efficient than IPv4, so it should mean faster and more reliable networks. And at the very least, it will mean that ISPs and IT departments worldwide will no longer need to deal with the uncertainty of not knowing when the addresses will run out.

accessible over the internet – if you're running a web server on a Raspberry Pi, for example, or a media server on a NAS appliance.


You can get around this using a clever feature that's built into TCP, which is the protocol that computers use to connect to one another (as distinct from IP, which defines the structure of the data they then exchange). Whenever one system initiates a TCP connection with another, it has to specify which virtual "port" it wants to communicate on. Conventionally, this gives a clue as to the purpose of the connection: for example, port 80 is normally used by web browser software wanting to connect to a web server, while port 401 is used for controlling and interacting with

an uninterruptible power supply.

You can use this information to point connection requests in the right direction. If you're running a web server on your Raspberry Pi, it will probably be listening on the default port 80. You can therefore go into your router's port-forwarding settings and configure it to forward all incoming connections on port 80 to the internal IP address of the Pi. The next time the router receives an incoming connection request on that port, it will direct it onwards to the Pi, and then pass the response back to the originating internet address.

You don't always have to set up port forwarding by hand. Some applications use a protocol called UPnP (Universal Plug and Play) to communicate with your router and automatically set up the required rules. Security experts have recommended disabling this feature on your router, however, because it can be used by malware to open up attacks from malicious servers. You might decide it's worth the risk: you should have good anti-malware defences in place regardless of whether you're using UPnP or not. ■

Your router provides internet services to all your devices.



SECRETS OF THE GOOGLE OFFICE APPS

TOP
SECRET

FOR
YOUR
EYES
ONLY

GOOGLE'S ONLINE WORD PROCESSOR AND SPREADSHEET
ARE MORE POWERFUL THAN YOU MAY REALISE;
NIK RAWLINSON EXPLORES THEIR HIDDEN FEATURES

Part of what's great about Google Docs is its familiarity. It looks and feels so similar to Microsoft Office, LibreOffice and other offline word processors that you can jump in and start working with very little culture shock. Similarly, Google Sheets looks enough like Excel that you can skip over any sort of training and start being productive right away.

However, if you think of Google's apps merely as free alternatives to the traditional desktop office suite, you'll miss out on some unique features, many of which take advantage of the online model. Here are some of our favourite features and tricks that will help you get the best from Google Docs and Sheets.

Pageless view

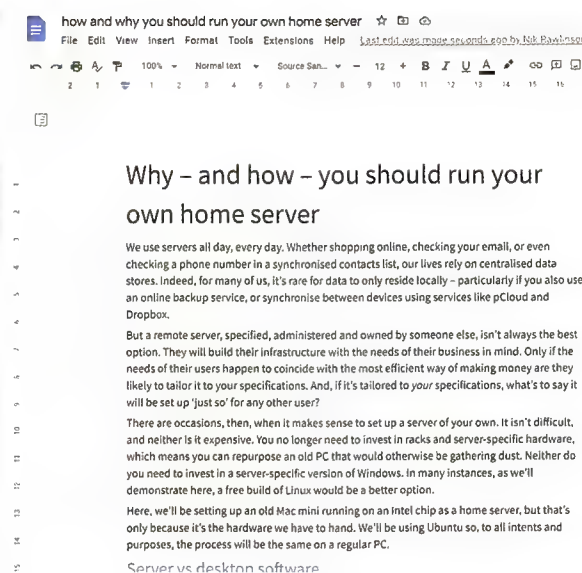
When word processors first became mainstream in the 1980s, it was a fair assumption that most things you typed would end up being printed out for sharing and consumption. That's certainly not true any more – so there's no need for many documents to be divided into pages.

To remove the virtual page breaks, open the File menu, click Page setup and switch to Pageless mode. Your page breaks will disappear, as will the page edges, giving your text room to breathe at the centre of the screen. You can also change the background colour to something other than white if you find that too glaring, or if you're working at night and want to reduce your blue light exposure.

Talk, don't type

Google Docs has a built-in voice-to-text processor, allowing you to dictate your documents as if you were talking to a virtual secretary. However, it only works if you're using Docs through Chrome, not any other browser, even if that browser is Chromium-based. You can activate it by pressing Ctrl-Shift-S, or by picking it from the Tools menu.

The translator doesn't just transcribe spoken words; it can handle punctuation, too, with full stop, comma, exclamation mark, question mark, new line and new paragraph all included in the range of commands it understands. When you've



finished writing, you can edit using your voice, too, using clear commands such as “select paragraph”, “select word”, “copy”, “paste”, “delete last word”, “insert page break” and so on. You'll find a full list of available commands in Google's help pages at tinyurl.com/APC515ghelp.


Type less

Voice typing is a great time saver, but if you prefer to keep your hands on the keys you can work more efficiently by setting up automatic substitution, to save retyping common phrases again and again. For example, if you sign off every document with your email address, you could set up Docs so you only need to type “my@” to insert your complete address.

To set up automatic substitution, click Tools

Google Docs can pull up relevant web links in a sidebar.

Save typing the same thing over and over with substitutions.



followed by Preferences and switch to the Substitutions tab. Type your shortcode (my@ in our example) into the empty box at the top of the left column, then put the full text in the corresponding box in the opposite column and click OK. Because this is set in the application's Preferences rather than the document itself, your new shortcode will work across all of your documents, new and old.

Do your research inside your document

Rather than switching back and forth between Docs and any web pages you need to refer to, you can access a lot of the research you might need inside Docs itself.

First, make sure you're not in the pageless view; then click the square speech bubble icon at the bottom of the screen (when you hover over it, an “Explore” label will appear). Google Docs will examine your document and pull up relevant web links in a tabbed sidebar that also includes options for Images (pulled from the web) and Drive (which contains any related documents in your Google Drive storage).

If the initial selection doesn't satisfy your needs, use the search box at the top of the sidebar to look for other relevant content.

Quickly format text

You no doubt know that Ctrl-B (bold), Ctrl-I (italic) and other conventional shortcuts work in Google Docs just as in

Word. You can also use Shift-Ctrl-comma and Shift-Control-full stop to make selected text incrementally smaller and larger, saving you a trip to the toolbar. Ctrl-comma and Ctrl-period alone invoke subscript and superscript respectively.

You can use similar shortcuts to reorganise text on the page. For example, you can select a section of text, then use Shift-Alt-up arrow to move it above the preceding paragraph, or Shift-Alt-down arrow to move it beneath the one below. You'll find a full list of keyboard shortcuts at tinyurl.com/APC515gshort – but be warned, it's long.

Add and edit images


Select Insert | Image from the Google Docs menu and you'll see options for uploading a new image, selecting one from Google Drive or using a picture in your Google Photos library. If you choose Google Drive or Google Photos, your online Drive and Photos storage will be opened in the sidebar; click the image (or images) you want to use, then click Insert.

Once the picture is on the page, you can perform basic formatting and editing functions without needing to open up an external image editor. Click the image to make handles appear; if you drag from any corner, the image will retain its original aspect ratio, while dragging from an edge lets you squeeze or stretch it. The uppermost handle has a stick, with which you can rotate the picture.

When an image is selected,

Upload and edit images without leaving Google Docs.

Bookmarks can draw attention to a specific part of a document.



you'll also notice that a new button – Image options – appears on the toolbar. Click this to set text wrapping options, change the image colour (to match your document theme or a corporate or club colour scheme, perhaps), and change its transparency, brightness and contrast. Right-click on an image, and you can add descriptive alt text, or crop the picture.

On a related note, there's an easy trick you can use to extract images from a Google Doc: open the File menu, then select the Download submenu and click "Web page (.html, zipped)". A ZIP file will then be downloaded to your computer, inside which you'll find a folder named "images" containing the images embedded within the document.

Link to a specific location

When a colleague first opens a document that you've shared with them, their view will start at the top. If you want

to draw their attention to a specific part, you can create a bookmark and share that instead.

To do so, position your cursor where you want them to start reading, then click Insert | Bookmark. A small tab icon now appears in the margin; click this to expand it, then click the rectangular Bookmark icon to copy a direct link to that location. You can send this link to anyone, as long as they have access rights to the document. You can grant these by clicking the Share button at the top of the screen and assigning rights to individual people, or making the document available to anyone with the link.

Insert an image

If you want to brand a spreadsheet with a company logo, include your signature at the bottom of the page or simply decorate it with pictures of your cat, you can easily incorporate an image hosted elsewhere online with the formula:

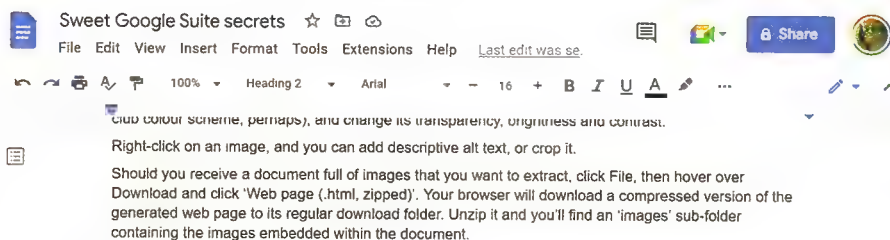
`=IMAGE("URL", mode)`

Note that the URL needs to be in quotation marks. The mode variable tells Sheets how to size and scale the image. If we wanted to include a photo at `example.com/pic.jpg`, we could use:

`=IMAGE("http://www.example.com/pic.jpg")`

This defaults to mode 1, which shrinks the image to fit within the cell while retaining its original proportions. Specifying mode two will stretch the image to fill the cell, like this:

`=IMAGE("http://www.example.com/pic.jpg", 2)`



Link to a specific location

Bookmark and, when the recipient opens it, they'll find themselves on line one. That's fine if you want them to read the whole thing. But, if you want to draw their attention to a specific part, create a bookmark and share that instead.

Sheets

Insert an image

If you want to brand a spreadsheet with a company logo, include your signature at the bottom of the page, or simply decorate it with pictures of your cat, you can easily incorporate an image hosted elsewhere online with the formula:

`=IMAGE("URL", mode, height, width)`

Mode 3 leaves the image at its original size, but hides any parts that extend beyond the edges of the cell, while mode 4 allows you to specify a custom height and width in pixels:

`=IMAGE("http://www.example.com/pic.jpg", 4, 480, 640)`

Insert a link

Inserting a hyperlink is a similarly simple exercise, with a formula that allows you to specify both the URL and the hyperlinked text that should appear in the cell.

`=HYPERLINK("https://history.gg/", "The history of Guernsey")`

The link text is optional: if you leave it out, the URL will be shown in the cell.

Collate data from multiple sheets

You can manually copy and paste data between Sheets documents, but if you want to copy a whole sheet from one document to another, you can do so in a few clicks. At the bottom of your Google Sheets document you'll see a tab for each sheet it contains; click the drop-down arrow on any tab, then select "Copy to" from the menu that appears. Select "Existing spreadsheet" and a file selector will appear; you can now pick another document to copy the sheet into.

Import live data from another spreadsheet

It's easy to collect sheets together as outlined above, but any copies you make aren't linked to the original data; if you update the source

document, the copied version doesn't change.

If you want your data to remain current in both documents, instead use the `IMPORTRANGE` function, which draws in live data from a separate document. The formula is:

`=IMPORTRANGE("spreadsheet_url", "range")`

The spreadsheet URL variable is just the address of the source spreadsheet in the browser, excluding the `/edit` on the end, and everything that appears after it. The range specifies the cells we want to include, such as `A1:B4`. Again, you'll need the quotation marks. If your source document has multiple sheets, you can specify a particular sheet by writing its name in front of the cell reference followed by an exclamation mark, just as in Excel. For example, if we wanted to import those cells from Sheet 2, we'd enter the range as `"sheet2!A1:B4"`.

When you first enter a formula such as this you may see an error, warning that you need to

Sheets' data validation lets you embed a drop-down list in a cell.

link the two sheets. Click the button to allow this and a live view of the source data will be imported into your main spreadsheet, running down and right from the selected cell. Any changes you make to the original cells will be immediately reflected in the second spreadsheet. (Note that the reverse is not true: if you overwrite any of the linked cells in the destination sheet, the whole block will be blanked out and you'll see an error message.)

Create drop-down lists



If you're using Sheets to design a form, you may want to make sure that any data entered conforms to a limited, pre-defined range. You can do this using Sheets' data validation function, which lets you embed a drop-down list within a cell.

To use it, click in the cell where the list should appear, open the Data menu and select "Data validation". Change the Criteria drop-down to "List of items", then type the entries for your list in the box to the right. Click Save; a down-pointing arrow will now appear at the end of the cell, which you can click to open the list.

You can create other list types just as easily. For example, if you want to offer values between 1 and 10, change Criteria to Number, and specify the range. Sheets can also handle dates, tickboxes and strings of text that either do or don't contain specific words.

Create smarter heatmaps

You doubtless know how conditional formatting uses colour to make data easy to understand at a glance. What you may not know is that you can customise the parameters to create visible colour thresholds. To do this, select the range of



You can create visible colour thresholds for your data.

cells containing your data, then open the Format menu and select Conditional formatting. Switch to the Colour scale tab on the overlay that appears, and choose a colour scheme by clicking the bars beneath Preview. You can now use the Minpoint, Midpoint and Maxpoint options to refine when each colour kicks in. Click Done to apply the colours to your data.

Use sparklines

Sparklines are tiny graphs that display inside a single cell, giving an instant visualisation of a series of data. You can create these in Google Sheets using the SPARKLINE function, with this syntax:

=SPARKLINE(reference, {options as an array})

The "reference" is the range of cells you want the graph to depict, such as A1:A4. The appearance of your sparklines can be configured by adding pairs of arguments – the first being the name of property you want to configure, and the second containing the value you want to set it to. For example:


=SPARKLINE(A1:A4, {"charttype", "bar"; "color1", "red"; "color2", "blue"})

There are four chart types to choose from: bar, column, line and "winloss", which shows a simple representation of whether successive values are positive or negative. You'll find full details of the various applicable options at tinypurl.com/APC515spark.

Prevent editing in selected cells

When you share a document from one of the Google apps, you can normally choose whether to make it read-only or to allow recipients to edit the content. With Sheets there's a third way: you can broadly permit editing, but prevent others from changing the contents of specific cells, unless they've been specifically authorised to do so.

To do this, select the range you want to lock, then right-click on the selection, select "View more cell actions" from the context menu, then click "Protect range" on the fly-out. Enter a description for the range, then click "Change permissions". You can now either allow specific people to edit, or show a warning when someone tries to make an amendment. In the



Sparklines let you display tiny graphs inside a single cell.

latter case, they can go ahead and make the change if they click OK, but they can't claim they weren't warned.

Get going quickly

When you're setting to work on a new document or spreadsheet, there's no need to step through multiple menus to get to a blank page: if you're already logged into your Google Account, just open a new browser window and visit doc.new (for a new word-processing document) or sheet.new (for a new spreadsheet).

Find hidden functions

If you're looking for a specific feature but can't remember where it is, hold down Alt and press /, then start typing a few letters from the function name. Matching options will drop down from the menu bar, and the selection will be filtered as you carry on typing. So, if you can't remember where Conditional Formatting is found in Sheets (it's in the Format menu), press Alt-/; type "cond" – and there it is, at the top of the list.

Share a PDF

You can turn any Google document or spreadsheet into a PDF in seconds: just open the File menu, then select Download | PDF document. An even quicker way is to click in the address bar, copy the URL and swap the "edit#" at the end for "export?format=pdf". You can then share this URL with a

correspondent; when they click it, they'll be prompted to view or save a freshly generated PDF, rather than the original editable document. Note that they'll need access to the original document for this to work, so make sure you've authorised them as a viewer or editor, or made the file accessible to anyone who has the link.

Create one-click templates

You can use a similar technique to create instant templates for new documents. Copy the URL of an existing file, but this time replace "edit#" with "copy". Now, when anyone accesses that URL, they'll be prompted to create a new copy in their own workspace, which they can go on to edit. Again, remember that whoever you share the link with will need access to the original (whether explicitly through being named, or implicitly because it's available to all).


Roll back your changes

If you go too far in an edit – or if someone else makes a mess of your work – you can always roll back. Open the File menu and select Version history; you can then browse the various editions of your file in the sidebar, in each of which new changes will be highlighted. Hovering over a change reveals who made it; you can revert the entire document to any previous state by simply clicking the "Restore this version" button. ■

PC BUILDER

System News

Why are new products hard to sell? Following Nvidia's woes, Mark Williams looks at why AMD is now having trouble.



Mark Williams
Mark is an IT professional with a strong interest in voiding warranties.

With so many CPU and GPU product launches in the last few months and more still to come into next year, we've never had so much choice when buying or building a new PC! With AM5's release though there have been a few gripes holding back consumers from making the jump to AMD's latest and greatest.

Firstly, there's the current cost of DDR5. With no dual DDR5/DDR4 support like Intel's current offerings, AMD's Ryzen 7000 DDR5 only path means a big wad of cash must be laid down on currently expensive RAM just to get your foot in the door, a tough sell when looking at the mid-range CPUs.

Second, we saw only expensive (but feature rich) X670/E

motherboards released at launch, with the B650/E boards released a few weeks later, which even then, were still very expensive. Richard from Aftershock PC told us "For instance, the B550 TUF Wifi II is going for \$279, while the B650 TUF Wifi is \$479. That is a hell of a jump in price for something that is not a premium 'X' series and is the 'lower-end' motherboard". AMD's mid-range chipset is simply too costly to be considered of value and is likely the main reason Intel 13th Gen is outselling AMD currently.

Third, AMD surprisingly since starting the core wars, is now losing on the core count front, not so much at the very top end but particularly with Ryzen 5 and 7 parts AMD is losing in

productivity to Intel where its E-cores are proving decisive.

Finally, there's AM4's long shadow. The Ryzen 7 5800X3D is such a good product that, as reported by mindfactory.de which publishes sales data on their sales in Germany for the DIY market, the 5800X3D by itself outsold the entirety of Intel's 12th and 13th Gen line-ups combined and with such close performance to Ryzen 7000 series parts at least in gaming, consumers are obviously not seeing the need to upgrade to an overpriced platform.

Better availability should lower motherboard prices over time and the launch of the Ryzen V-cache parts early next year may be the tipping point for AMD to turn the tables once again on Intel. ■

SHOP TALK

How are AMD's 7000 series CPU sales going against Intel's 13th Gen parts?

Jaimie, Leader Computers:

"We have seen the market swing back to team blue for new performance builds in the past weeks.

What is interesting and maybe affecting this somewhat is the poor availability of motherboards from the big three vendors for the AMD platform. I think we will get a better perspective in a few months as supply of the motherboards stabilises. The situation is improving now but the AM5 launch was probably four weeks early to get maximum traction."

Andre, JW Computers:

"Intel is on top of the game at the moment because it has all the right upgrades in all the right places.

With double the number of efficiency cores over the 12th Gen, currently better backwards compatibility and value for money. At the end of it all, customers just want something reliable and reasonably priced with the option to upgrade in the future. Intel seems to be moving a lot easier as it can work out to be very expensive very fast if you want the AM5 Zen 4 line-up.

Market Watch

A sampling of PC systems available this month



CPL ONLINE INFINITY i5

\$1,499 | tinyurl.com/APC515CPL

This value-oriented system does many things right but stumbles in a few areas. The system overall with its 650W power supply, case, motherboard and RAM are all great for what's on offer here, the trouble comes when looking at the cooling, GPU and storage. Using a CPU from a tray might save CPL a few dollars but it means no heatsink is included so one must be added to the bill, raising the cost.

The sole storage device is a 500GB SSD, that's all, no secondary HDD here for more storage. Lastly, the GPU was never good value at launch and is still overpriced with its 30 percent performance deficit to a similarly priced RX 6600. An ok system but there is better value to be found elsewhere.

CPU: Intel i5-11400; **Cooler:** Antec A30; **Motherboard:** Gigabyte GA-B560M-D3H; **Graphics:** Gigabyte GeForce RTX 3050 Eagle 8GB; **Memory:** Crucial Ballistix RGB 16GB DDR4 3200MHz CL16; **Storage:** Crucial P2 500GB M.2 NVMe SSD; **Power Supply:** Antec VP650P-PLUS 650W; **Case:** Antec VSK10.



BPC KRAKEN POWER SYBER

\$5,888 | tinyurl.com/APC515BPC

Sporting AMD's latest 12-core CPU, this rig is aimed at gamers that require and demand the best from their PCs.

With big ticket items like 360mm AIO, 32GB DDR5 RAM, 2TB SSD, 1000W PSU and top-end X670 motherboard, there's almost nothing that isn't the best you can buy. Oddly it's the CPU that's not the absolute best, sitting one rung below the Ryzen 9 7950X. Will that matter for games? Nope, it's in a strange middle ground between the productivity happy 7950X and the gamer ideal 7700X. So long as you do a bit of both, you won't be wasting your money here. Overall, an extremely fast and capable rig for gamers and prosumers alike.

CPU: AMD Ryzen 9 7900X; **Cooler:** MSI MAG Coreliquid 360R AIO; **Motherboard:** Gigabyte X670 Aorus Elite AX; **Graphics:** Gigabyte GeForce RTX 4090 Gaming OC 24G; **Memory:** Kingston Fury Beast RGB 32GB 6000MT/s DDR5 CL36; **Storage:** Kingston 2TB KC3000 PCIe 4.0 NVMe M.2 SSD; **Power Supply:** Gigabyte GP-UD1000GM 1000W; **Case:** MSI MPG Velox 100P Airflow.



AFTERSHOCK PC MAINFRAME - EXTREME

\$5,999 | tinyurl.com/APC515AFT

Coming in a little dearer than the BPC rig, this system starts off on a similar footing with the 13700K which matches the 7900X. With 32GB of RAM and equally high-end motherboard as well, things start to diverge from there; there's 1TB more storage but 1TB less of it on SSD, the PSU its smaller capacity, and the big one – the GPU is only an RTX 3080 as specced.

The lesser specs are a necessary fact of having to cover the costs of the included gorgeous custom hard line open water-cooling loop. If you're thinking of buying this system, we'd recommend considering an RTX 4000 GPU when purchasing, as taking this loop apart to upgrade later won't be an easy task.

CPU: Intel Core i7-13700KF; **Cooler:** Mainframe EK Custom Open Loop water cooling; **Motherboard:** Asus Prime Z790-P WIFI-CSM DDR5; **Graphics:** Asus RTX 3080 ROG Strix Gaming OC V2 – 10GB; **Memory:** 32GB Kingston Fury Beast Black 5200MHz DDR5; **Storage:** 1TB Gen4 Samsung M.2. NVMe, 2TB Seagate Barracuda HDD; **Power Supply:** 850W Inwin; **Case:** Fractal Design Torrent.

JW COMPUTERS GMR APOLLO 6600

\$1,395 | tinyurl.com/APC515JWC

If you're looking to maximise the performance-per-dollar spent on a system, it's hard to look past this rig from JW.

Sure, it uses a quite old CPU now with the Ryzen 5 3600 but so long as you play games that stress the GPU more than the CPU this will do just fine.

Speaking of the GPU, the bang-for-buck king RX 6600 resides within and will handle 1080p gaming perfectly and even 1440p with some game settings tweaks. Mated with a B550 motherboard there's plenty of room for CPU upgrades later and with a beefy enough PSU, GPU upgrades within reason will also be a non-issue. The included 2TB HDD is the cherry on top relative to the CPL system making this an easy budget recommendation.

CPU: AMD Ryzen 5 3600; **Cooler:** OEM; **Motherboard:** Asus Prime B550M-K; **Graphics:** Asus Radeon RX 6600; **Memory:** Kingston Fury Beast RGB 16GB DDR4 3600MHz; **Storage:** 500GB NVMe M.2 SSD, 2TB HDD; **Power Supply:** Asus TUF Gaming 750W; **Case:** Coolermaster MasterBox MB520 ARGB.



Blueprints

Value- and performance-driven hypothetical builds

BUDGET

A perfect balance between price and performance.

With the release of the latest generation of processors and GPUs, we were expecting a fall in the price of components this time around, especially in the budget department. Yet, here we are, scanning through all the sites, only to find that our CPU and GPUs, on the whole, are much the same price as last month.

Maybe all the new launches raised our expectations for these builds and we're yet to feel the effects of this trickling down into the budget market. Or perhaps we should just be grateful that prices of components aren't still skyrocketing, as they were this time last year. After all, at this price point, you can still grab a fairly powerful machine capable of 1080p gaming.

Due to prices barely fluctuating, both builds remain almost the same, which admittedly, doesn't make for the most exciting blueprint we've ever done. However, one other component that has changed is the RAM. As pointed out by one of our eagle-eyed readers, it isn't necessary to use 3,600MHz RAM for our budget machines. These chips can only run at speeds of 3,200MHz, so we have added 3,200MHz sticks into both budget machines instead. If you're using faster memory then it will still probably be supported by the motherboard, just make sure you enable XMP in the BIOS.

On the AMD machine, we swapped out our previous PNY XLR8 Gaming Epic-X RGB sticks for the equivalent 3,200MHz variant, though the price turned out to be basically unchanged. As for the Intel build, out goes the PNY XLR8 Gaming Epic-X RGB sticks and in comes Corsair's Vengeance RGB Pro SL running at 3,200MHz.



"Maybe all the new launches raised our expectations for these builds and we're yet to feel the effects of this trickling down into the budget market."

AMD INGREDIENTS

Part		Price
Case	Corsair 4000D Airflow	\$120
PSU	550W Cooler Master MWE Gold 80+ Gold	\$120
Mobo	MSI B550M PRO-VDH WiFi AM4	\$149
CPU	AMD Ryzen 5 5600X	\$299
GPU	MSI Radeon RX 6650 XT Mech 2X OC 8GB	\$640
RAM	16GB (2x 8GB) PNY XLR8 Gaming EPIC-X RGB @ 3200MHz NEW	\$109
SSD	500GB Corsair MP600 PCIe 4.0 M.2 SSD	\$160
HDD	4TB Seagate SkyHawk Surveillance 7200	\$129
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	\$60
Total		\$1,786

INTEL INGREDIENTS

Part		Price
Case	Corsair 4000D Airflow	\$120
PSU	550W Asus ROG Strix 80+ Gold	\$140
Mobo	ASRock B660 Pro RS LGA 1700	\$179
CPU	Intel Core i5-12600	\$359
GPU	Zotac GeForce RTX 3060 Gaming Twin Edge OC 12GB	\$580
RAM	16GB (2x 8GB) Corsair Vengeance RGB Pro SL @ 3200MHz NEW	\$118
SSD	500GB MSI Spatium M450 PCIe 4.0 M.2 SSD	\$89
HDD	4TB Seagate SkyHawk Surveillance 7200	\$129
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	\$60
Total		\$1,774



"The AMD mid-range build sports the new Ryzen 5 7600X CPU, which features six cores, 12 threads, and a maximum boost turbo frequency of 5.3GHz."

MID-RANGE

A game-ready machine that can also handle demanding work.

Our mid-range builds are where we start to see the introduction of next-generation components. At the time of writing, we are still being introduced to the latest parts from both Intel and AMD, and in the previous issue, we brought in new CPUs, which will remain at the helm of their relevant ships.

The AMD mid-range build sports the new Ryzen 5 7600X CPU, which features six cores, 12 threads, and a maximum boost turbo frequency of 5.3GHz. We also have one eye on futureproofing this build, as the RYZen 5 7600X features support for

DDR5 RAM up to 5,200MHz, the new AM5 socket, and Zen 4 architecture – overall, it's a highly capable processor and at a reasonable price point too. Of course, by making that change, we had to bring in a compatible motherboard – the Asus Prime B650M-A AX Micro ATX AM5. This also has support for DDR5 RAM at 5,200MHz and so matches perfectly with our 32GB of Corsair Vengeance DDR5 CL40 that runs at the same speed.

These prices stayed the same, so we didn't see the need to swap them out. However, we did replace the previous

Nzxt C750 PSU with the same rated Corsair model, giving us a fully modular form factor over the semi-modular one we had previously.

We made a small typo last month, sorry to say we left the important 'Ti' off the Gigabyte GeForce RTX 3070 Ti Gaming OC 8GB. That's corrected now, and apologies for that error.

Unfortunately, there were no further changes on the other side of the fence for our mid-range Intel build. Other than a few changes in prices we mentioned, the PSU was also the only change made here too.

AMD INGREDIENTS

Part		Price
Case	Nzxt H7 Flow	\$219
PSU	750W Corsair RM750 80+ Gold NEW	\$189
Mobo	Asus Prime B650M-A AX Micro ATX AM5	\$399
CPU	AMD Ryzen 5 7600X	\$515
Cooler	Cooler Master Masterliquid ML240L RGB V2	\$95
GPU	Asrock Radeon RX 6800 XT Phantom Gaming D OC 16GB	\$899
RAM	32GB (2x 16GB) Corsair Vengeance DDR5 CL40 @ 5200MHz	\$269
SSD	1TB Crucial P5 Plus NVME M.2 PCIe 4.0	\$179
HDD	4TB Seagate SkyHawk Surveillance 7200	\$129
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	\$60
Total		\$2,953

INTEL INGREDIENTS

Part		Price
Case	Nzxt H7 Flow	\$219
PSU	750W Gigabyte P750GM 80+ Gold NEW	\$189
Mobo	MSI PRO Z690-A ATX LGA 1700	\$279
CPU	Intel Core i5-13600K NEW	\$549
Cooler	Cooler Master Masterliquid ML240L RGB V2	\$95
GPU	Gigabyte GeForce RTX 3070 Ti Gaming OC 8GB	\$1,099
RAM	32GB (2x 16GB) Corsair Vengeance DDR5 CL40 @ 5200MHz	\$269
SSD	1TB Crucial P5 Plus NVME M.2 PCIe 4.0	\$179
HDD	4TB Seagate SkyHawk Surveillance 7200	\$129
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	\$60
Total		\$3,067



"We can finally introduce an RTX 4090 GPU to the mix, and boy, have we been waiting patiently for this!"

TURBO

The maximum PC.

With a larger budget and an arsenal including the latest performance-focused PC components to play with, our Turbo machine blueprints have been far more exciting than the other two. That said, most of the big changes were made in the previous issue, when we introduced Intel's new powerhouse, the Intel Core i9-13900K, and AMD's equivalent Ryzen 9 7900X processors. On the AMD side, we made a total of five changes to the machine last time around and, unfortunately, that means it's not quite as exciting this month.

However, the Radeon 6950 XT GPU is falling in price so we've added in a MSI Radeon RX 6950 XT Gaming X Trio 16G.

We are still awaiting AMD's latest GPUs to compete against the RTX 4090 found on the Intel Turbo machine, and although the Asrock OC Formula card is a mighty powerful GPU that is more than capable of 4K gaming, it doesn't stop us from being eager to get our hands on the latest AMD GPUs.

Although we made plenty of changes to the Intel machine last time around, there was one thing we couldn't change back then. Thankfully, we are now in the position to make that pretty exciting swap. Yes, we can finally introduce an RTX 4090 GPU to the mix, and boy, have we been waiting patiently for this! We've specifically chosen the

Asus GeForce RTX 4090 TUF Gaming 24GB as it was the cheapest available of the bunch. We say "cheapest", but in real terms, it's anything but cheap. Costing more than both our budget machines combined, the GPU comes in at almost \$3,000. That's a steep price, but it's also a steep performance upgrade with plenty of futureproofing too, considering the lifespan of the previous card and the longevity the RTX 3090 Ti still has today.

There's quite some disparity between the two Turbo builds, but this should change when we get the new AMD graphics cards in.

AMD INGREDIENTS

Part		Price
Case	Phanteks Enthoo Pro 2 Tempered Glass	\$220
PSU	1000W Nzxt C1000 80+ Gold NEW	\$249
Mobo	Asrock X670E PG Lightning ATX AM5	\$539
CPU	AMD Ryzen 9 7900X	\$849
Cooler	Cooler Master MasterLiquid ML360R RGB 66.7 CFM 360mm	\$249
GPU	MSI Radeon RX 6950 XT Gaming X Trio 16G NEW	\$1,899
RAM	64GB (2x 32GB) Corsair Vengeance DDR5 CL40 @5200MHz	\$509
SSD	1TB WD Black SN850X M.2 PCIe 4.0	\$249
HDD	6TB WD Blue 5400 HDD	\$189
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	\$60
Total		\$5,012

INTEL INGREDIENTS


Part		Price
Case	Phanteks Enthoo Pro 2 Tempered Glass	\$220
PSU	1000W EVGA SuperNova 1000 GT 80+ Gold	\$280
Mobo	Asus TUF Gaming Z690-Plus WIFI DDR5 LGA 1700	\$389
CPU	Intel Core i9-13900K	\$1,099
Cooler	Nzxt Kraken X73 73.11 CFM 360mm	\$239
GPU	Asus GeForce RTX 4090 TUF Gaming 24GB NEW	\$2,959
RAM	64GB (2x 32GB) Corsair Vengeance DDR5 CL40 @5200MHz	\$509
SSD	1TB WD Black SN850X M.2 PCIe 4.0	\$249
HDD	6TB WD Blue 5400 HDD	\$189
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	\$60
Total		\$6,193

QUICK TIPS

The problem solvers

The APC team field problems across the spectrum of devices and software.

Learn a new trick or fix you can use.



Move old emails offline to free up online space.

Confusion over mesh

I've got a couple of dead spots in my home network. The obvious solution would appear to be a mesh network, but there's so much choice – even from the same manufacturer sometimes. What sort of things should I be looking for when shopping for a mesh network?

Joel Dawson

APC responds: There are many things to consider when looking at mesh networks. If you're looking to futureproof your mesh network, then make sure it supports WiFi 6 as a minimum, and WiFi 6E if you have lots of devices where competing demands for network bandwidth could exacerbate your problems.

WiFi 6 adds an additional 160MHz channel to the 5GHz spectrum, so those devices capable of using WiFi 6 (802.11ax) are moved on to a separate channel away from other 802.11ac (WiFi 5) devices to boost performance and capacity. The major difference between WiFi 6 and WiFi 6E is that the 'E' stands for extended, extending the WiFi 6 channel into the 6GHz band to get an extra 1,200Mbps bandwidth for increased performance and capacity.

WiFi 6E charges a hefty premium (around \$150) and with many people still primarily using older WiFi 5 devices, there's little

to be gained right now by choosing it over WiFi 6. Other factors come into play, namely the number of bands provided by the router and the maximum advertised speeds.

You'll see routers with listings such as AC2400 or AX4000. The 'AC' or 'AX' identifies them as either WiFi 5 or WiFi 6/6E, while the secondary figure is a theoretical maximum speed offered across all available bands. For example, an AC1750 dual-band router might offer up to 450Mbps on its 2.4GHz channel, and 1,300Mbps on its 5GHz channel.

A tri-band WiFi 6 router can deliver improved performance across three channels, for example, the Asus ZenWiFi XT9 AX7800 router divides its capacity to offer up to 574Mbps over its 2.4GHz channel, 2,402Mbps on the WiFi 5 5GHz channel, and 4,804Mbps on the WiFi 6 5GHz channel. Proving that not all WiFi 6E routers offer superior top speeds, Amazon's new eero Pro 6E only supports up to 5,400Mbps across its three channels: 600Mbps over 2.4GHz, then 2,400Mbps on each of the other channels.

These speeds are all theoretical, of course, which is why it's important to also see what range is offered as, ultimately, your choice of router should come down to whether it delivers decent

performance over an extended range as opposed to how theoretically fast it is when you're stood next to the router itself. In recent tests, for example, we found the ZenWiFi XT9 offered faster speeds at range than the eero Pro 6E.

Other things to look out for include a dedicated 'backhaul' channel. Some mesh devices, including some tri-band WiFi 5 systems, allow you to reserve a channel exclusively to channel traffic between the router and satellites to keep the rest of the bandwidth free for your networking needs. Support for this backhaul often helps extend the range too, allowing you to push each satellite that bit further away from the main router. Also, check to see that WPA3 is supported – most support it, but beware of those that don't include a WPA2 fallback option for older devices when it's enabled, such as Google's Nest WiFi.

Finally, before investing in a new mesh network, check to see if your current router supports mesh. Both Asus and Synology bake mesh support in to their latest routers, so you may already find yourself halfway to a solution. And if you're only looking to extend your network to one far-flung location way out of range of your router, consider a HomePlug option instead.

Free up email storage

I don't pay for Google One storage and am running out of space on my account, which I use primarily to store emails. Should I move to a provider that has more free storage, such as Yahoo Mail, use an app to strip attachments from emails, or pay for extra storage?

Jennifer McDowell

APC responds: Your options are many and varied, and none involve having to switch email account provider. Of course, you could pay for extra storage – Google One's Basic tier costs \$1.99 a month for 100GB of free storage, which should be ample. But why

pay when you can simply make better use of your existing storage?

One option is to archive older messages offline, freeing up storage without having to make any awkward decisions over what to keep and what to get rid of. APC has been using a free tool called MailStore Home to keep his unruly Google email in check for years.

Alternatively, you could switch to accessing your Gmail through Thunderbird (www.thunderbird.net). Once your account is connected, create a filter to move all messages after a specified date into a local folder, which would then delete the online copy to free up space. Click 'Manage message filters', then click New... and set up a match filter including 'Date' and 'is before', then choose to 'Move Message to' a local folder. Alternatively, you could set up a filter using the Size (kB) parameter to move only those messages containing attachments over a certain file size.

Thunderbird also allows you to 'detach' individual attachments from messages. Save the file to your hard drive and it is deleted while the original email remains online. Sadly, there's currently no way to perform this action in bulk – one such add-on, AttachmentExtractor, doesn't work with the latest version of Thunderbird. Consequently, APC recommends the archiving route as your best option.

Noisy fan fix

One of the fans on my desktop PC has started to make an incredibly loud whine and I'm not sure what to do to resolve it.

Wiley S. Sanders

APC responds: The first thing to do is identify the problematic fan. Is it your CPU, case, GPU, or maybe even your power supply? Once you've identified the culprit, take a closer look. Is the fan clogged with dust or some other obstruction? Try cleaning it using a can of compressed air or a soft brush, then see if that has any effect on noise levels. Also, make sure there's plenty of room for air to circulate around both fan and your PC.

If your fan is rattling or grinding, then it's likely to be coming to the end of its life, so replace it if possible. If looking to replace a case fan, check with the case manufacturer or measure it carefully (including any screw holes) before replacing it.

CPU fans are also straightforward to replace – make sure you choose the right fan for your CPU model as well as its make, and if you're looking for a bigger fan, make sure there's enough space in your case to accommodate its larger footprint. Once your new fan arrives, run your PC for a while to heat up the CPU

and soften the thermal paste connecting it to the heatsink.

Then, carefully remove the old fan before cleaning the CPU die to remove any leftover thermal paste using 91 percent proof alcohol (72 percent would be adequate) and lint-free paper. Some new fans come with thermal paste pre-applied, while others may require you to apply this yourself, so check online videos to ensure you use the right amount for maximum heat transference.

GPU fans can be trickier to replace. If your card is still under warranty, see if the manufacturer will offer a free replacement.

Otherwise, you should first verify the fans are removable – many, but not all, are. They may be attached via clips or small screws to a passive heat sink (some may be obscured by stickers), and you should also make a note of the

power connector and where it goes.

Even if the fan can be removed, you should source a replacement before proceeding. You will need to search for 'fan' plus your video card make and model to see what's out there.

Finally, the trickiest fan to replace is that on a PSU. You'll need to open up the PSU to identify what the current fan is and whether it's replaceable, then repeat the process when the new fan arrives to actually swap them out. Given the potential for giving yourself a nasty (or even lethal) electric shock, APC is wary of recommending such an approach – instead, replace the PSU itself.

File Explorer freeze

Something odd is happening in

Windows: I can drag, drop and double-click files as normal, but as soon as I attempt to right-click a file or folder File Explorer becomes unresponsive. How do I fix it?

William Erickson

APC responds: This is usually linked to a shell extension, added by a third-party program into the context menu to provide access to that application. Identifying it is the hardest part, so install ShellExView, free (www.nirsoft.net/utils/shexview.html).

Once installed, launch the program and choose 'Options > Filter by Extension Type'. Click Context Menu so it's the only highlighted entry and hit OK. Now, click the Company column header to the right of Product Name to sort the entries, making it easier to group the

non-Microsoft entries together – all are marked in pink.

Hold Shift as you click on the first and last pink entry to select them all, right-clicking the selection and choosing 'Disable Selected Items', then clicking Yes when prompted. Now choose Options > Restart Explorer (again, click Yes when prompted). Try right-clicking now – if the problem goes away, you know it's one of the entries you've disabled.

Now re-enable each entry one at a time, then restart Explorer, until the problem reappears. Once you've identified the culprit, leave its shell extension disabled, remove the program or troubleshoot further using its website. ■


"If your fan is rattling or grinding, then it's likely to be coming to the end of its life, so replace it if possible"



Asus's ZenWiFi XT9 offers a great base for a mesh system.

Stop your photos and videos decaying

Don't let 'data rot' distort and destroy your precious memories. Robert Irvine explains how to safeguard your image and video files against damage and loss in quality.



Unlike printed photos, which can be damaged by creases, rips and stains, we tend to assume that our digital snaps and videos are indestructible. But some Google Photos users recently found that pictures they'd uploaded to the service a few years ago showed signs of ageing like physical photos, including cracks, discolouration and blurring that looked like water damage.

Digital files themselves don't degrade over time, but – as proved by the Google Photos flaw – other factors can make their data distorted and unreadable. Here, we explain how you can

prevent digital decay to preserve the quality of your photos and videos.

"Digital files themselves don't degrade over time, but – as proved by the Google Photos flaw – other factors can make their data distorted and unreadable."



Save photos in a lossless format

Although digital photos don't gradually decay, every time you edit and save an image file it loses some of its

quality. This is especially true when using the JPEG format, which automatically compresses data, sacrificing some of the original bytes to reduce the file size.

Many image editors allow you to save photos at the best JPEG quality to keep data loss to a minimum – for example, IrfanView (www.snipca.com).

Darktable lets you work with uncompressed RAW images that never degrade.



IrfanView saves JPEGs with minimal loss of quality, but other formats are superior.

com/43599) has a JPEG 'Save quality' slider you can set to Good, and an additional option to 'Try to save with original JPG quality' (see screenshot right). But the nature of the format means your photos will still be compressed, and the more edits you make the more bytes you'll lose.

You can avoid this problem by saving photos in the lossless format TIFF. This enables image files to store lots of information and means their quality won't be affected when you edit them. Most digital cameras let you take photos in TIFF format, though the resulting files will be much larger than JPEGs. Your phone's camera probably uses JPEG or HEIF (if it's an iPhone) format, but you can convert its images to TIFFs in a desktop image editor.

If file size is a problem, because you don't have much storage or want to share your pictures by email or online, try saving in PNG format. Unlike TIFF, this uses compression to reduce file size, but it's also lossless so the quality of your images won't degrade. However, because PNG files don't store as much data as TIFFs, and don't support CMYK colour – which is essential for high-resolution, 'true colour' prints – the format is best suited for web images.

You can 'upscale' photos in

JPEG, TIFF or PNG format, to improve their resolution without affecting their quality, using the free tool Reshade (www.snipca.com/43596). This excellent program can also refocus blurry pictures and reduce photo defects.

Capture photos in RAW format

For the best possible quality, and to safeguard images against digital decay, capture your photos in RAW format.

This stores an unaltered copy of the data generated by your camera, so no data is lost. You can then keep copies of the original RAW files, while saving edited versions of photos as TIFFs and PNGs. To shoot in RAW quality on your phone, install an app such as Open Camera for Android (free from www.snipca.com/43592) or Halide Mark II for iOS (\$4.49 a month from www.snipca.com/43593).

RAW's main drawbacks are that it creates very large files – usually at least 20MB – that need to be processed using compatible software such as Darktable (www.snipca.com/43591, see screenshot below left) or a plug-in in your usual image editor, such as RAW FileType for Paint.NET (www.snipca.com/43590). This allows you to adjust attributes such as contrast, saturation and white

balance, whereas other image formats process this data automatically.

Edit and convert videos without losing quality

As with photos, the quality of videos can degrade each time you edit and save them, particularly when using a 'lossy' (as opposed to lossless) compressed format such as MP4. When you convert a video to another format or reduce its size, the file is 're-encoded', which results in data being lost and the quality deteriorating.

One solution is to install the free open-source tool LosslessCut (www.snipca.com/43594). As its name suggests, this lets you edit video – and audio – files, to remove unwanted segments, and convert them to different formats, without re-encoding them and losing quality. The program supports all the main video formats, including MP4, MOV, MKV and H264 – the latter offers the best balance between quality, file size and compatibility.

The only catch is that you'll need 7-Zip (www.snipca.com/43611), or another tool that supports 7z files, to open LosslessCut's download file – avoid its Microsoft Store app, which costs ~\$30. Once the program is installed, load your video, click the point in the timeline where you want it to begin and select 'Start current segment', then click the desired


end point and choose 'End current segment'. Click 'Export' (see screenshot below) and select a format to save your video in without any reduction in quality.

Store multiple backup copies of your files

The main cause of data rot comes not from within the photo, video or document, but the storage media it's saved on. Storing a precious file on a SD card may be convenient, but if that card becomes corrupted, then you'll lose the memory forever.

Although CDs and DVDs should theoretically last 50 years, they're susceptible to scratches, mould and disc rot, and cheap discs often warp and become unreadable. Internal and external hard drives may seem secure and offer plenty of space for your files, but they may eventually fall prey to bad sectors or malware that corrupts their data. Even good old USB sticks can be mislaid, or lose the small electrical charges required to store data in their flash memory.


Because no storage format is infallible, it's a good idea to have a minimum of three copies of your most valuable photos and videos – one original and at least two backups stored on separate devices. Storage locations should include a web-based copy in a service such as Google Drive or OneDrive, which gives you extra security in case your home is burgled or suffers fire or flood damage.



Google Photos might seem like a logical place to back up media files, but be aware that – in addition to its recent glitch – it will store your photos and videos at a "slightly reduced quality", unless you go into its Settings and select 'Original quality' (see screenshot above). Your uploads will be larger and use more of your Google account's storage limit, but at least their data won't be altered. ■

LosslessCut lets you edit and save videos without affecting the picture quality.

Choose 'Original quality' in Google Photos to save full quality photos and videos.




FIX DAMAGED PHOTOS AND VIDEOS FOR FREE

Most tools that claim to repair corrupted photos and videos will charge you, but they may provide a free, watermarked preview of the fixed file, so you can at least confirm they work. Such is the case with Stellar Repair for Photo (\$46 from www.snipca.com/43597) and Stellar Repair for Video (\$45 from www.snipca.com/43598). However, unless you have a lot of damaged files, try fixing the data rot yourself before you spend money.

Changing the name or format of an unopenable or distorted image file can sometimes repair the damage. Right-click the photo, choose Rename and enter a new

name, then try opening it using a different image editor – we recommend the versatile IrfanView (www.snipca.com/43599). If it still looks corrupted, click File then 'Save as' and save the picture in another format, such as PNG, to see if that helps. To change the format of files that can't be opened, upload them to the online conversion tool Online Convert (www.online-convert.com), which processes images of up to 25MB for free.

VLC Media Player (www.snipca.com/43600) can often convert corrupted videos to a working file format. Open its Media menu, select Convert/



Save and click the Add button. Choose the broken video file, then click the Convert/Save button and choose a codec from the Profile dropdown

menu (see screenshot). Select a folder to save the file to, then click Start to convert and hopefully repair the item.

MAC OS

Master the Clock app

At long last, the Clock app comes to desktop and notebook Macs.



IT WILL TAKE
5 minutes

YOU WILL LEARN
How to see world times, set alarms, and use the stopwatch and timer

YOU'LL NEED
macOS 13


The Clock app does more than tell the time. It offers world time clocks so you can see what time it is elsewhere; you can set alarms that trigger at a time of your choosing; and it has a stopwatch with a digital or analogue display and a lap timer, as well as a countdown timer that lets you know when a certain period of time has elapsed.

The Clock app has been a popular feature of the iPhone and iPad for many years, but with macOS Ventura, it finally comes to the Mac too. It doesn't sync with iCloud, so your world clocks and alarms don't carry across from your mobile devices, but you can always add them separately if you wish. In this guide, we show you how to get the most from your Mac's

Clock app, with a guide to its basic features and a few advanced capabilities that you might otherwise miss. Here's a tip right now. The app icon, whether in the Dock, a Finder folder or the App Launcher, accurately shows the current time, all day, every day.


Ian Osborne

HOW TO USE THE CLOCK APP'S WORLD CLOCK



1 ADD A WORLD CLOCK

Open the Clock App, and click on the World Clock tab at the top. Your current location's time zone is already shown. To add a new world clock, click the '+' icon (top-right). Type the name of a city in the Search field.



2 ADD ITS TIME ZONE


When you've found your city, click on it to add its time zone to the Clock app. The light and dark sections of the world map show the areas currently in daytime or night-time, as do the white or black clock faces.



3 DELETE A CLOCK


To delete a world clock from the Clock app, move the pointer over its clock face and a 'X' appears in the top-left corner of its window. Click on this cross and that time zone disappears. You can add it again later.

HOW TO SET AN ALARM



1 SET AN ALARM


Click the Alarm tab, then click the '+' icon in the top-right corner. Click on the hours then minutes, and type the time you want to set. You can then set which days you want the alarm to trigger, give it a label and more.



2 CHOOSE A SOUND

Click the Sound menu to choose a sound for the alarm. Then click the Save button to set the alarm. To edit an alarm, click on it and make changes. To delete it afterwards, hover over it and click the 'X' icon in the top left.

HOW TO USE THE TIMER



1 SETTING THE TIMER

The countdown timer is useful for signalling when a certain time has expired, for example, when baking a cake. Click the Timer tab, and click the hours, minutes and seconds, typing the values you want to use.

HOW TO USE THE STOPWATCH



1 TIME SOMETHING

Click the Stopwatch tab. Click the green Start button to start timing, then click Lap to record a time without stopping the clock, or Stop to stop it altogether. The Reset button wipes it zero, so you can start again.




2 CHOOSE OTHER FUNCTIONS

In the Clock app's View menu, you can switch between a digital or an analogue display for the Stopwatch. They do the same job. You can choose between the Clock app's various functions in this menu too.



2 SETTING THE TIMER

The countdown timer is useful for signalling when a certain time has expired, for example, when baking a cake. Click the Timer tab, and click the hours, minutes and seconds, typing the values you want to use.



STOP PLAYING

An overlooked feature of the Timer

The Timer section of the Clock App has an extra feature of which you might not be aware. At the foot of the 'When timer ends' menu is an option called Stop Playing. Choose this option and whatever audio you're playing on your Mac stops when the countdown timer reaches zero. It's great for falling asleep to relaxing music or the radio, and having it turn off automatically after a certain amount of time. At the time of writing, the Stop Playing option doesn't work on the Mac, though the option is there, so it will probably be fixed soon. ■

"It's great for falling asleep to relaxing music or the radio, and having it turn off automatically after a certain amount of time."

LINUX

Track down files and folders instantly

'Work smarter, not harder' advocate Nick Peers reveals how to use a tool to transform the way you search for files and folders on your PC.



QUICK TIP!

Most virtualisers can still boot floppy images. In VirtualBox, go to the Storage tab within the settings for a virtual machine and add a new controller of type floppy. Then, add a floppy image in the same way that you do an ISO image.

Having a day when you're running out of RAM and wish you had a faster processor? This issue, we're looking at operating systems that can boot almost instantly, run from a floppy disk and require as little as 8MB of RAM to boot. You might think we're delving into systems of the past, but these are currently maintained operating systems.

Specifically, we're looking at MenuetOS, an operating system coded in pure assembly language for speed and a frugal memory footprint, and a system forked from it called KolibriOS. Obviously, systems like these are stripped down for speed and simplicity, and we don't expect them to take over from operating systems such as Windows and Linux on the typical desktop. However, they can do quite a lot and they are fascinating to examine and play around with. Like a lot of alternative OSes, they offer a glimpse of the direction that operating systems might have taken if the computer industry had evolved differently.

MenuetOS

MenuetOS is a free download from the official website (www.menuetos.net). It started life as an open source 32-bit operating system. However, it then transitioned to being a 64-bit OS and became closed source, although it's still free to download and use.

Initially, we started with the smallest 64-bit build on the site, the floppy image, partly because we were curious to see what an operating system on a 1.44MB storage medium could offer, and we booted it using the VirtualBox virtualiser. The first thing that hit us was the boot speed, as it takes only a few seconds before you are presented with a full desktop. Amazingly, the desktop is populated with icons that link to many applications, with more in the application launcher.

Exploring what's on offer, it's obvious that this is a highly customised operating system

```
SHELL 0.8.3
KolibriOS v0.7.7.0, Kernel SUN-rev.: 9872
Type 'help' for help
# help
31 commands available:
about      alias      cd
clear      cp         mv
ren        date      echo
exit       free      help
history    kfetch    kill
kill       ls         lsmode
mkdir      more      ps
pwd        reboot    rm
rmdir      shutdown sleep
touch      uptime   ver
waitfor
```

The terminal offered by KolibriOS is simple, but it works. It has a command history but no tab completion.

with a unique set of applications. There is a file manager, a text editor, a number of simple games and even a web browser, along with some other internet utilities. This is in addition to other utilities such as a video player and a fairly complete suite of GUI system monitors and setup utilities. Just to assure you that you've not misread, we are still talking about the floppy edition of MenuetOS! Although this is astonishing on a technical level, it's not quite miraculous. The applications themselves tend to be very simple. For example, the GUI text editor has only the basic facilities, with no advanced features such as syntax highlighting.

The custom web browser can display the MenuetOS website and works surprisingly well on a few other sites, such as the search page of Google. However, performance isn't great, which is probably due to architectural limitations of the OS.

MenuetOS can read and write to FAT16 and FAT32 partitions. It seems a shame that there is no actual installer program supplied. If you've got the skills, instructions exist to use another operating system to partition a hard drive, copy system files and set up a bootloader such as GRUB. The developers are obviously happy for most users to boot from a floppy drive or flash drive to explore what MenuetOS can do.

KolibriOS

KolibriOS (<https://kolibrios.org>) was forked from the 32-bit version of MenuetOS in 2004. The two operating systems are similar and much of what can be said about one applies to the other. It does look like KolibriOS will remain 32-bit only for the foreseeable future, meaning it can only be used as a learning platform for x86 rather than AMD64 assembly language. A point in its favour, KolibriOS is fully open source and released under the GPLv2 licence.

Assessing it in terms of the distro, we feel it offers a more complete user experience than MenuetOS, and it comes with more applications. Unfortunately, MenuetOS and KolibriOS have diverged to the point that they are no longer binary compatible with one another. The GUI is more refined, but the biggest advance in that department is that it makes extensive use of right-click invoked pop-up menus, a frustrating omission in MenuetOS. A trip to the forum confirmed that there is no way of increasing the font size that all applications follow. This is the GUI's biggest let-down.

The applications

As was the case with MenuetOS, to explore the software library is to enter a parallel universe of non-standard, custom tools. CEdit is a more powerful text editor than the default MenuetOS one. It offers syntax highlighting and is a fairly

There's quite a big collection of simple games. If you're in the mood to steer tetriminos into tight spaces or land a little spaceship, you should find something to while away a few minutes here. There are also some ports of larger games such as the shareware editions of Doom, Wolfenstein and Quake. There are also ports of emulators such as DOSBox and ZSNES. The range of graphical demos and games that come with KolibriOS show that it can offer good performance for software-based rendering, but as is the case with HaikuOS, there is a lack of

So, how do these unusual, minimalist systems work?

What we say here about the technical underpinnings of these two related operating systems can be taken to apply to both MenuetOS and KolibriOS, unless stated otherwise.

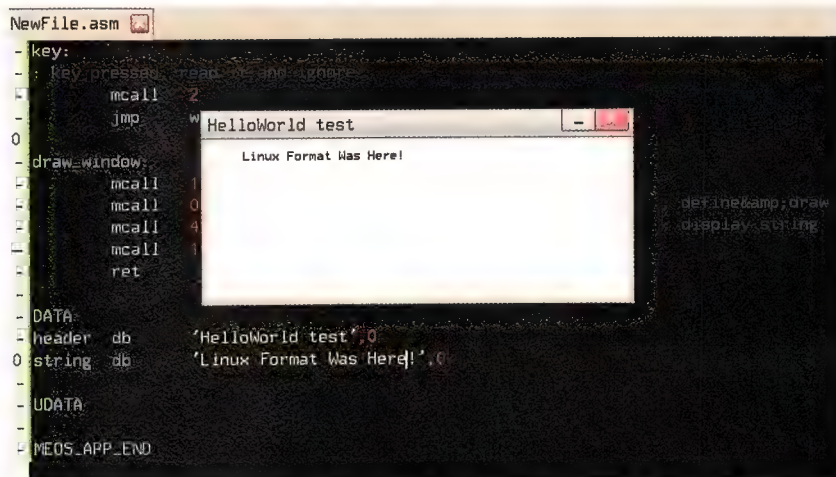
The reason that more mainstream operating systems don't use a microkernel is that it imposes a significant performance penalty. It would probably be a difficult sell to most desktop users if they were offered a more secure and stable operating system that ran more slowly, particularly when the stability of most mainstream operating systems is deemed to be 'good enough'.



The MenuetOS kernel does offer pre-emptive multitasking – the kernel controls the multitasking of applications by allocating CPU time and protected memory. To understand this, consider a simple program that prints text to the screen. Under an OS that provides preemptive multitasking (such as Linux or Amiga OS), that program wouldn't have to know anything about the underlying operating system.

MenuetOS has a different look from KolibriOS and both systems have configuration potential but, disappointingly, no control over the default GUI font.

The MenuetOS-derived operating systems tout themselves as good platforms for no-nonsense development using assembly language. But what is that? The earliest digital computers were programmed in machine code, the language of a CPU. On early computers, rows of switches loaded a set of values into Assembly is an easier way to program machine code that uses mnemonics, small commands, rather than numbers. **mov ax, #60** moves the value of 60 into the 'ax' register. **add ax, bx** adds the contents of the ax and bx registers and deposits the result in the ax register. This is how most games were created until the mid-'90s. Each assembly language is specific to a type of CPU. It produces small code that is quick to run, but working with it has been likened to building a ship in a bottle.



We built a Hello World program using the tools provided and altered it slightly.



QUICK TIP!

A lot of alternative operating systems use the VESA standard to determine the available screen modes. You can add extra screen modes (such as 1,366x768 resolution with 32-bit colour) to VirtualBox with the following command: `VBoxManage setextradata "[name of VM]" "CustomVideo Model1" "1366x768x32"`

has to process messages from the OS, and if it stops for a moment or crashes, the rest of the system grinds to a halt. It's also impossible for such a system to ensure the smooth operation of background processes. Much CPU time is wasted while the foreground application is waiting to complete its processing.

When an application can launch multiple processes that operate simultaneously, it's called multithreading. MenuetOS doesn't provide multithreading facilities, but it can launch system processes that execute simultaneously. This is one of the reasons why it can boot to a desktop so quickly and general operations seem so fast; it's because one slow system component doesn't force the rest of the system to also wait.

This is a point that raised some controversy in the Linux world a few years ago as most of the major distros switched to a system for service management



Kolibri Far is one of the file managers that comes with KolibriOS and will suit fans of text-based interfaces.

called Systemd and away from SysVinit. The essence of Systemd, when it comes to booting, is that it launches many services simultaneously during boot. This means that a service that is slow to start up doesn't block the rest of the system. The fact that resources

are saturated during boot means the boot is far faster. The downside is that it's a more complex startup procedure and initially it was more bug-prone.


In short, MenuetOS can run multiple applications, but like AROS and RISC OS, it can't launch multiple threads within an application or service. In addition to causing performance problems during the operation of complex applications, this limits the potential for porting large, modern apps such as web browsers from other operating systems, as they tend to rely on multiple threads that are managed by the OS. In all fairness to MenuetOS, it's one of the fastest-booting graphical OSes we've seen. However, that's because the developers have deliberately kept the design simple. Although it doesn't support multithreading, MenuetOS does support multiple CPU cores.

POSIX

Operating systems such as Linux are mostly POSIX compliant. POSIX (Portable Operating System Interface) is a set of standards to make sure that an OS offers certain programming interfaces, so that code can be reasonably portable. MenuetOS doesn't aim for POSIX compatibility at all, and this is reflected in the available tools and applications that it ships with.

Adding POSIX support to MenuetOS could alter the character of the system, as it means adding a huge amount of complexity. As it stands, MenuetOS has only a handful of important configuration files, and it is also capable of operating with a tiny memory overhead.

Compare this to AROS, which we



GET AROUND KOLIBRIOS

1. TASK SWITCHER

This works the same way as it does on other operating systems.

2. SOFTWARE

There are lots of juicy applications and games to play around with in the KolibriOS floppy disk edition.

3. FILE BROWSERS

KolibriOS comes with a couple of different options for file management.

4. CONFIGURATION

Both operating systems can be configured using GUI tools.

5. SYSTEM MONITORING

Each operating system comes with the type of system monitoring tools you'd expect.

6. MORE APPS

Some of the highlights of the application library can be located on the backdrop.

looked at last month. AROS isn't fully POSIX compliant, but it does offer many of the required facilities. For this reason, AROS comes complete with ports of most of the GNU coreutils package that provides the familiar commands such as ls, cat and rm, and ports of other Unix-type operating system tools. Overall, as you're exploring MenuetOS, it has a rather alien feel to it as most of the tools it comes with have been created from scratch.

Developing applications

MenuetOS has its own windowing system that isn't compatible with any other system, another factor that limits the potential for ports from other systems. In theory, it would be possible to port one of the standard, cross-platform GUI toolkits, such as GTK, but no such port exists at the moment.

Internally, MenuetOS's graphical toolkit works in a similar way to other graphical toolkits. It's event-based, meaning that the application waits until it receives a message from the system, such as one telling the application it has received keyboard input or needs to redraw its window. Programmers who have worked with toolkits such as Qt or GTK should be able to grasp how it works, but the biggest problem is that MenuetOS is orientated towards assembly language programming, which most modern programmers are

less familiar with. MenuetOS applications can be created in languages other than assembly language, such as C, but nearly all the example code uses assembly language. Development in languages other than assembly language has to be carried out on another operating system and transferred over to MenuetOS for testing.

All in all, it would take a big commitment to make a GUI application for these systems.

Mini marvels

It's interesting to contemplate that what these operating systems offer would once have constituted a fully-featured operating system. The applications that KolibriOS and MenuetOS offer are minimalist. The biggest let-down of the GUI in either case is that there isn't a central control to increase the font size that all applications and GUI components respect. In terms of what the kernel offers, in both cases, it's impressive that they support multiple CPU cores and pre-emptive multitasking. However, it's a shame that they don't offer multithreading support, partly for performance reasons and partly because it would make porting mainstream applications more of a possibility.

As targets for original and ported app development, both OSes have a couple of things against them. Adding POSIX



compatibility would greatly increase the bloat and add performance overheads and complexity. Lack of ports of standard GUI toolkits like GTK and extensive support for languages other than assembly also makes porting large applications difficult.


The developers are probably reluctant to change the entire character of these operating systems, which can boot almost instantly, boot from a floppy and run under 8MB of RAM. We can't say that application development is impossible, because there are probably over a hundred games, apps and utilities available for both systems. It's a shame the two systems are now incompatible. If they could remerge them, it would offer a larger single developer and user community for this fascinating style of minimalist OS. ■

Most of the operating system configuration is carried out graphically, and there is a full complement of applets for this.

HARDWARE COMPATIBILITY

KolibriOS and MenuetOS have similar hardware support. KolibriOS has a slightly lower minimum CPU requirement, as the currently maintained build is a 32-bit one, and a Pentium class processor is required. For MenuetOS, pretty much any 64-bit CPU from a Core 2 Duo onwards should be suitable.

Both operating systems tout their ability to run on systems with 8MB of RAM, but we found that booting within a VM with only 8MB didn't work properly. It's probably technically possible to boot under those circumstances if you were prepared to go in and edit some configuration files and strip the system down a bit. We caused some crashes in KolibriOS by bringing the desktop up with only 16MB of RAM. So, we recommend 128MB (an eighth of a gigabyte) as a reasonable minimum for operating the desktop and loading some of the native applications, because this did work. Neither OS has the ability to



When booting KolibriOS, you can specify the screen mode to use and the ramdisk backup to load.

swap RAM out to a hard disk.

USB is supported for common peripherals, such as the keyboard and the mouse. Beyond the basic VESA

support, there are also special drivers for ATI graphics cards and various common soundcards.

Arduino to Pi GPIO comms using Python

Like a snake charmer, Les Pounder uses his skills to persuade an Arduino to talk to a Raspberry Pi via Python.



YOU'LL NEED THIS

- Raspberry Pi
- Latest Pi OS
- Arduino Uno (or compatible)
- Breadboard
- 10K potentiometer
- LED
- 330-Ohm resistor (orange-orange-brown-gold)
- 5x M2M jumper wires
- Code: <https://github.com/lesp/LXF-pyFirmata>

The Arduino and the Raspberry Pi both have their strengths and weaknesses, but what if we could harness both of them to produce projects? Firmata is a firmware that enables our Raspberry Pi (or another Linux computer) to talk to an Arduino and use it as an extension of the GPIO. In this tutorial, we shall learn how to configure an Arduino that reacts to commands sent via pyFirmata, a Python version of Firmata. This project would be ideal for those old Arduino clone boards that many makers have in their bits box, pressing them into service to provide extra GPIO and analogue inputs that are missing on our favourite single-board computer. Using pyFirmata, we can easily add extra features to our projects, including those that just need a few more GPIO options.

Preparing the Arduino Uno

Our Arduino Uno requires a firmware to be flashed upon it. This firmware, Firmata, acts as a translator between our Python code and what the Arduino can understand. You need to download and install the Arduino Legacy IDE (1.8 series) from www.arduino.cc/en/software. At the time of writing, Arduino IDE 2.0 has just been released, but it only provides 64-bit installation files.

Raspberry Pi OS is 32-bit. Download the installation archive and extract it to a directory in your home directory. Open a terminal and change directory so you are inside the Arduino directory.

Run the install command using `sudo`:

```
$ sudo ./install.sh
```

When complete, we need to add our user to a group called `dialout`. This is a group for users who can access serial ports. On a Raspberry Pi, the default user is typically `pi`, but change this to your username:

```
$ sudo usermod -a -G dialout pi
```

Reboot your Raspberry Pi to complete installation. Our next task is to connect and test our Arduino, so it can receive the firmware. Connect your Arduino to the Raspberry Pi via a USB cable, then open the Arduino IDE. Go to Tools > Board and set your board. In our case it was an Arduino Uno. Then go to Tools > Port and set the port for your board, this is typically `/dev/ttyACM0`. Next, load the Blink sketch from File > Examples > 01 Basics and click on Upload (arrow pointing to the right) to compile and upload the sketch to

the board. You should see the Arduino's LED blink every second. This proves our connection works and we can now upload the Firmata firmware to the Arduino. Click on Load File > Examples > Firmata > StandardFirmata and upload to your Arduino. When done, close the Arduino IDE.

To talk to our Arduino, we need to install a Python3 module, `pyFirmata2`, using the `pip3` package manager. `Pip` is normally installed on a Raspberry Pi, but for other distros it may

"This project would be ideal for those old Arduino clone boards that many makers have in their bits box."

not, so we've included a line to install it:

```
$ sudo apt update
$ sudo apt install python3-pip
$ pip3 install pyfirmata2
```

Close the terminal and go to the Raspberry Pi main menu, Programming > Thonny.

Our first project with `pyFirmata2` is to control an LED, connected to GPIO pin 6. This LED will blink 10 times, then we will set the brightness of the LED using PWM (Pulse Width Modulation). See the circuit diagram in the download for details on how to wire the project.

Import the `pyFirmata2` and time modules:

```
import pyfirmata2
import time
```

Create an object, `PORT`, that stores the serial port for our Arduino, then create another object, `board`, to make a connection between our code and the board.

```
PORT = pyfirmata2.Arduino.AUTODETECT
```


```
board = pyfirmata2.Arduino(PORT)
```

Now use a for loop that iterates 10 times:

```
for i in range(10):
```

Inside the for loop, turn digital GPIO pin 6 on (1) then wait for 0.2 seconds before turning the pin off (0) and wait a further 0.2

The Arduino and Raspberry Pi have their strengths and weaknesses. Using Firmata, we can get the best from both boards.



seconds:

```
board.digital[6].write(1)
time.sleep(0.2)
board.digital[6].write(0)
time.sleep(0.2)
```

Out of the for loop, we reach the section controlling the PWM LED. Create an object, `pwm`, that we use to configure GPIO pin 6 as a digital pin (`d`) which uses PWM (`p`):

```
pwm = board.get_pin('d:6:p')
```

Use a variable, `v`, to store the user's choice of number between 0 and 100. This is then divided by 100 to produce a float value between 0 and 1...

```
v = float(input("PWM duty cycle from 0 to 100: ")) / 100.0
```

Set the LED brightness by writing the PWM duty cycle to the `pwm` object:

```
pwm.write(v)
```

Pause for five seconds, then turn the LED off and close the connection to the Arduino.

```
time.sleep(5)
pwm.write(0)
board.exit()
```

Save the code as `pyfirmata-blink.py` and click on Run. It takes a few seconds before anything happens, but you should see the LED blink, and then the Python shell asks for your input.

Analogue to digital

The Arduino has many analogue inputs that we can use to measure voltage changes in a circuit. For the next project, we will use a 10K potentiometer to control the brightness of an LED. See the circuit diagram in the download for details on how to wire the project.

From the `pyFirmata2` module, import the `Arduino` class and then import the time module:

```
from pyfirmata2 import Arduino
import time
```

Create a global variable, `pwm`, that we can reference inside and out of functions:

```
global pwm
```


Auto-detect the port and Arduino board, then set GPIO pin 6 to use PWM:

```
PORT = Arduino.AUTODETECT
board = Arduino(PORT)
pwm = board.get_pin('d:6:p')
```

Create a class, `Analog_To_PWM`, which will contain the code for the project:

```
class Analog_To_PWM:
```

Create a function inside the class to set the sampling rate to 10Hz and to connect with our Arduino:



The analogue circuit is made up of a circuit for the LED, connected to GPIO 6, and a circuit for the 10K potentiometer, connected to A0.

```
def __init__(self):
    self.samplingRate = 10
    self.board = Arduino(PORT)
    Create a start function that will use analogue pin 0 (A0) to sample data and then set it to report data to the pyFirmata:
```

```
def start(self):
    self.board.analog[0].register_callback(self.myPrintCallback)
    self.board.samplingOn(1000 / self.samplingRate)
    self.board.analog[0].enable_reporting()
```

Another function, `myPrintCallback`, reads the data from the analogue input and prints it to the Python shell. We also use the value to write the PWM duty cycle to the LED. The returned values are rounded to one decimal place for easy reference:

```
def myPrintCallback(self, data):
    print("A0 reads,%f" %
```

```
(round(data,1)))
```

```
pwm.write(round(data,1))
```

The final function ceases reading the analogue pin:

```
def stop(self):
    self.board.samplingOff()
    self.board.exit()
```

Create an instance of the code, ready for use:

```
analog_to_pwm = Analog_To_PWM()
```

Run the code using a try, except handler to try to start the code; if it fails, the code stops:

```
try:
    analog_to_pwm.start()
except:
    analog_to_pwm.stop()
```

Save the code as `pyfirmata-analog-to-pwm.py` and click Run to start. Wait for a few moments and you will see the PWM value printed to the Python shell; turn the potentiometer to change the brightness of the LED. Click Stop to end the code. ■



GENIUS TIP!

The Arduino's GPIO uses a 5V logic, which makes it incompatible with the Raspberry Pi's 3.3V GPIO. Don't directly connect the two as you risk damaging the Pi's GPIO. If you must connect them, use a voltage divider or a logic level converter.

PYFIRMATA AND ALTERNATIVES

The `pyFirmata` (<https://github.com/tino/pyFirmata>) module was created by Tino de Bruijn over a decade ago as a means to use Python with the Firmata protocol (<http://firmata.org>). The version we use in the tutorial is a fork, `pyFirmata2` (<https://github.com/berndporr/pyFirmata2>) by Bernd Porr, which in recent years has seen updates and bug fixes. Porr's description says: "PyFirmata2 turns your Arduino into a data acquisition card controlled by Python" and has "up to 1kHz precise sampling at the analogue ports for digital filtering". This makes `pyFirmata2` a very interesting choice if you

want to capture data using an Arduino.

But what about other boards? MicroPython and CircuitPython are both receiving lots of love from the maker and microcontroller community. Largely thanks to boards such as the ESP32, STM32 and Raspberry Pi Pico. To interact with these, we can use Brian Pugh's Belay (<https://github.com/BrianPugh/belay>), which interacts with Micro/Circuit Python boards via REPL (Read, Eval, Print, Loop) over serial and Wi-Fi. Using Belay, we can send commands to a board, read pin status etc, all using pure Python.

GAMES

Call of Duty: Modern Warfare II

Proves Infinity Ward is still the CoD studio to beat.

■ \$109.95 | PC, PS4/5, XBS/X | callofduty.com

apc


HIGHLY
RECOMMENDED

Modern Warfare II's greatest strength is that it doesn't have a gimmick. And that it doesn't need one to stand out. All *Modern Warfare II* needs to be is a great shooter, and here it delivers. It is the *Call of Duty* experience, deconstructed, reconsidered and obsessively optimised over 15 uninterrupted years of iteration.

I know when I'm playing a *Black Ops* game because it has offbeat arcade modes and the latest iteration of Nuketown. I know I'm playing a *Modern Warfare* game because the shooting is so finely tuned that I want to take my Gunsmith build out to dinner. Infinity Ward has officially mastered the instant gratification of aiming down sights and deleting targets.

That goes a long way, though some of *Modern Warfare II's* biggest swings like revamped progression and a genre-bending campaign are less elegant. It also suffers from an obtuse UI, fits of crashes, weird bugs and the unexplained absence of basic features like stat tracking. But altogether, it's a good year for *Call of Duty*.

Ballistics are once again physically simulated, but you wouldn't know it from playing on regular 6v6 maps. On arena maps with short sightlines, guns behave like lasers that instantly tag targets. Jump into a 64-player Ground War match and the same guns



accommodate to firefights spanning entire city blocks, requiring players to lead shots and account for bullet drop. It's kind of wild how flexible and seamless the system is.

As a counterbalance to *Modern Warfare II's* wicked-quick ballistics and time-to-kill, Infinity Ward has pumped the brakes on operators themselves. In terms of movement, this is by far the slowest CoD in recent memory. Established traversal techniques like slide cancelling and bunnyhopping, much to the dismay of CoD's loudest fans, have been deliberately abolished.

By gluing our boots more firmly to the ground, *Call of Duty* has finally created the space to be a more methodical, dare I say tactical, FPS. For the first time in years, it is actually a better idea to tread lightly, hold angles, and listen for footsteps than full-sprint down every straight and bunnyhop around corners.



The campaign's 17-mission run is a healthy mix of familiar corridor shootouts and digressions into gimmicky experiments. You won't find Treyarch-style multiple endings or sidequests, but there are a few truly memorable missions.

Modern Warfare II is, happily, exactly what I wanted from another *Call of Duty*. As the stewards of *Call of Duty* filling the gap left by no new game coming in 2023, and the upcoming *Warzone 2.0*, Activision is putting its full weight behind Infinity Ward as a lead studio. *Modern Warfare II* is a good argument that they're the right folks for the job.

Sets a new bar for *Call of Duty* all over again.
Shame about the crashes though.

Morgan Park

★★★★☆

apc

HIGHLY
RECOMMENDED

VAMPIRE SURVIVORS

A roguelike that can win over the roguelike-avoidant.

■ \$7.49 | PC, XB1/S/X, Mac, Linux | store.steampowered.com

Vampire Survivors is built around a horde mode in simple, sprawling maps. You pick a character from a selection of Belmont-alikes then try to outlast a horde of ghoulies and ghosties, growing in power and number over a 30-minute timer. The game's first great curveball is how it handles shooting. Each weapon is mechanically unique, with different AoEs, fire rates, and damage profiles. Instead of directly targeting enemies, the weapons have a timed firing pattern influenced by you and your enemies' positioning.

With six weapons maximally upgraded and at least a few of them 'evolved' through item combos, your character lets out a constant torrent of projectiles pushing back an endless wave of foes. It practically plays itself at this point, but that's part of the fun. The closest thing I can compare it to is when you whip up a deathball of units in an RTS or more tactical RPG. It pokes that same lizard brain pleasure centre for me as when I drag to select a doomstack of Battlecruisers in *StarCraft*, or a partyful of mages with Melf's Minute Meteors in *Baldur's Gate*, click on an enemy, and

watch the sparks fly.

I certainly enjoy *Vampire Survivors* on my desktop, but I think it really sings on Steam Deck (or, failing that, a thin and light laptop). I got most of my gameplay for this review in on my couch or out on the porch. Untethered from a desktop-grade CPU, *Vampire* does start to chug a bit in those late run steamroll sequences, but the game doesn't really demand twitch reflexes and I honestly kind of like the effect – it feels like my Deck is straining under the weight of all those meteors. Creator Poncle plans to transition it to a new, hopefully more

stable, engine by year's end.

Vampire Survivors is a killer little game, a fun roguelike that absolutely hooked me. It'll also only kill your free time, not your wallet.

Vampire Survivors would be worth a look even if it wasn't the price of a latte. It also works well on low-spec PCs.

Ted Litchfield

★★★★☆

"With six weapons maximally upgraded and at least a few of them 'evolved' through item combos, your character lets out a constant torrent of projectiles"

KEVIN COSTNER'S WATERWORLD

We need to talk...

■ Free | PC | bit.ly/WaterworldGame

Slowly but surely, all of *The Simpsons'* fake videogames are being made a reality. We've had *Lee Carlo's Putting Challenge* – which Marge bought Bart instead of the gory *Bonestorm*, remember – and now here's *Kevin Costner's Waterworld*, the machine that Milhouse played with quarters at the Springfield arcade.



What does *Waterworld* look like as an arcade game? *The Simpsons* never imagined what lay beyond that initial sidescrolling scene, but developer Macaw45 has gone back to the movie for inspiration. Mainly they've added thugs that you have to gradually punch to death. You'll die to the first enemy – as you acclimatise to the timing of your and your enemies' punches – but it's simply a matter of learning that timing, of knowing when to move in, punch, and then quickly back out.

Of course, you die every time the game changes, but this is just the game reinforcing its punchline: of Milhouse handing over a small fortune to one of the most cruel arcade games ever made. It's a really frustrating game, for sure. But, like Milhouse, you'll want to see it through.

Often tedious and frustrating, but seeing a *Simpsons* joke being brought to life is a constant joy.

Tom Sykes

★★★★☆



SWAPWOOD QUEST

Battling with blocks.

■ Free | PC | bit.ly/SwapwoodQuest

Created for the latest Ludum Dare game jam – which had the theme “every ten seconds” – *Swapwood Quest* feels like a fully formed match-three game. You don’t control the falling blocks themselves. Instead, you shift them about after the fact, swapping any two blocks on the board horizontally.

At first, you’re merely trying to keep the blocks from reaching the top of the screen, but it’s not long before enemies are introduced. In each stage you’re given a creature to defeat in pitched combat. You do this by matching sword blocks and shield icons that will spare one of your HP.

Every time you match these combat blocks, a token is placed in a little window at the bottom of the screen, ready to be acted upon when the time is right. That time is every ten seconds, when both the player and enemy combat tokens are executed simultaneously. As yours and the enemy’s tokens are weighed against each other, either or both will suffer damage, and then the ten-second countdown will begin again.

Flaws are easy to forgive in a game this cute.
Tom Sykes

★★★★★



apc

HIGHLY
RECOMMENDED



THE CASE OF THE GOLDEN IDOL

Solve a string of bizarre murders.

■ \$25.95 | PC | thegoldenidol.com

A woman lies dead on a sofa, her mouth open in shock. A few members of the dinner party she was attending stare in horror at the scene while others level angry accusations at each other. There’s only one person smiling happily about the woman’s grim fate, and that’s me, the detective. I’ve got another intriguing murder to solve in *The Case of the Golden Idol*, one of the most novel and challenging mystery games I’ve played in years.

I begin my investigation, clicking on anything and everything, starting with the guests. They’re frozen in time, stuck in the moment the victim perished, so I’m free to go through their pockets with my mouse pointer to peer at their coins, rings and monogrammed handkerchiefs. A waste bin contains discarded correspondence, a dinner table holds unfinished meals, and most interestingly, a pouch dangles on a rope outside a window as if someone was trying to conceal something from the rest of the household.

Over the next hour I slowly piece together the identities of each person and learn who masterminded the

killing, who carried it out and for what reason, not to mention a few interesting plot twists along the way. This murder, and likewise this entire game, is a darkly brilliant tale told in discrete clues, and it makes me feel pretty clever for piecing them together and figuring it out.

There were a couple cases that didn’t entirely come together for me. In one I couldn’t fully figure out the motive, and I wound up occasionally brute-forcing the final few names or details of a case, trying every possible combination with the empty slots. But just because a few pieces didn’t fall into place didn’t stop the overall

story from reaching a satisfying end.

Paired with the satisfaction that comes from solving a complicated crime and moving on to the next, there’s a long and fascinating story running through these murders, all centred around the Golden Idol itself.

A challenging and inventive set of mysteries tied together with a brilliant storyline.

Christopher Livingston

★★★★★

"There’s only one person smiling happily about the woman’s grim fate, and that’s me, the detective."



BONELAB

A bold, weird VR physics sandbox with a story.

■ \$49.99 | PC, Meta Quest 2 | store.steampowered.com

Half-Life Alyx never truly wowed me, because Boneworks and its simulationist approach to VR already had its hooks in me. Physical avatars on 3D legs and no invisible walls to stop me. *Bonelab* expands on that experimental design. While tough for newcomers (moving around on virtual legs can be demanding until you get a feel for it) *Bonelab* is one of the most exciting experiences I've had in VR, despite the inevitable physics jank.

Where *Boneworks* was a traditional linear action game, *Bonelab* is closer to *Garry's Mod* for VR, interwoven with a five to six hour story. Set in an alternate '90s where VR took off early, the story has you breaking out of a conservative metaverse to become a free agent.

Mostly through VR running, gunning, brawling and puzzle-solving, but with a deeply creative and open endgame. *Bonelab's* early areas teach fundamentals. Movement, basic physics puzzle solving and combat. While not overly realistic in their handling, the guns feel satisfying to fire, with better sounds and effects than *Boneworks*, although melee has seen the most improvement.

Melee in *Boneworks* never quite felt

right. Larger weapons felt unwieldy and slow. In *Bonelab*, swinging a sword or club feels far easier. Blunt attacks land satisfyingly and blades pierce virtual non-flesh, requiring force to dislodge. Bare-handed fighting is even better, letting you grapple and even headbutt.

A warren of secret areas and Easter eggs, rewarding new tools, gadgets and

NPCs to summon is half of how *Bonelab* sells the fantasy of breaking free of VR's limits. The other half is an invite to the official modding Discord channel. Despite the mod SDK not yet supporting new AI or custom items, the community has already done that and more. *Bonelab* is still also in flux. Since launch, there have been updates to both mod tools and the


"Where Boneworks was a traditional linear action game, *Bonelab* is closer to *Garry's Mod* for VR, interwoven with a five to six hour story."

game. It's anyone's guess where it'll end up. I'm just excited to see how *Bonelab* – or VR itself – will change in the coming months.

Pushing boundaries. Strange, playful and committed to creativity but perhaps not for newcomers to VR.

Dominic Tarason

★★★★★



SALLY'S SCRIBBLES

Great art takes seconds.

■ Free | PC, Mac, Linux | bit.ly/SallysScribbles

It must be tough, being an artist in demand. Every day, clients insist you draw a cat or a smiley face or some broccoli or a line graph – and to draw it in only ten seconds. That's the premise for *Sally's Scribbles*, a game about making it as an artist – or not.

At least the requests are simple enough to draw. You can just about manage a cat's face or a stick man in ten seconds, though things get harder when you're asked to draw things in specific colours. The trick is that, for the most part, a vague squiggle will actually do instead – clients have been happy when I've rebelled and drawn a cross instead of a circle.

Does this revelation undermine the premise, or is it simply because it's art, and people will accept my 'abstract' version of an apple? Personally – after feeling like I'd seen behind the curtain, and potentially ruined the game for myself – I came to enjoy the freedom of drawing however I wanted.

I enjoyed the first few, frantic days, before the harsh realities of the artist life hit home.

Tom Sykes

★★★★★





Quick button presses can help mitigate incoming damage.

GAME CHANGER

Costume Quest

Tyler Colp is enjoying the spooky goings-on.

■ Release: 2010 Developer: Double Fine Publisher: In-house Link: bit.ly/3fcCQas

Costume Quest is a tradition I wish I had. Every time October sneaks back up on me, I get the urge to replay one of my favourite little RPGs. It's a small-scale spooky adventure full of witty writing and trick-or-treat nostalgia – and it's not a bad turn-based RPG either. But every year Halloween comes and goes and I don't manage to load it back up. This year, I finally did it.

The truth is that *Costume Quest* is timeless. You don't need to play it in the autumn. It exudes so much Halloween energy that you could play it on a summer's day, close your eyes, and hear the wolf howls and bat flutters that make up the ideal spooky atmosphere. Double Fine made the ultimate game for anyone that spent their childhood Halloween evenings roaming a cul-de-sac with a plastic bag full of candy. *Costume Quest* is that vibe, frozen in time and covered in a light RPG adventure wrapper.

Brush away a little bit of the 2010 jank – the game's inconsistent 60fps can be unwieldy along with some gritty textures – and you have a game

that is legitimately wholesome. The threat is positively Saturday morning cartoon: lazy, bug-eyed aliens stealing the candy from every home. When you knock on the door, you might get a happy adult or a Grubbin, the latter of which kickstarts the game's battles. And as you progress, you pick up discarded cardboard and cloth that help you craft new costumes.

Nothing about *Costume Quest* is complicated, and that's why it's easy to return to. When so many big games want to be everything all at once, it's pleasing to play a game that limits its ambition to do one thing incredibly well. And despite that small size, it captures the scale of what it's like to be a kid running along the sidewalks and streets, running into teachers and friends on the way to your next house.

Where other games continue to invent new ways to turn casual actions into game mechanics, *Costume Quest* puts all its effort simply into simulating a mood that you'll enjoy.

"The game's bright, cartoon-ish characters could use a little help"

Candy hunt

Costume Quest's combat remains perfectly engaging and relaxing. I still can't force myself to skip the intro to each fight after having seen every unique one years ago. Your chosen sibling and their goofy allies prepare for fights by transforming into fully-realised versions of their dinky costumes. I'll never forget the first moment Wren leaves the cardboard helmet and wings behind and explodes into a blue Gundam-like mech. It's a brilliant way to explain why battles are happening in the first place. *Costume Quest* is keenly aware of how Halloween is a chance to let our imaginations fill in the blanks.

The battles themselves have just enough to chew on that it doesn't get overly repetitive by the end of the game's eight-hour length. You take turns attacking and defending with the enemy Grubbins until they get knocked out and give up their candy. The only complication is that you can make the relatively easy fights go a little faster with quick time



The neighbourhood is full of minigames and people to talk to.

events. If you nail a button prompt mid-attack, it'll do more damage. Every costume has different versions if it and many of them are randomised. The penalty for missing them isn't severe, but if you get into the rhythm you can slide through battles nearly untouched.


Boss battles are more spectacle than challenge. The fact that I barely remembered the game had climactic battles with a giant grim reaper, a witch and a bulldozer, says everything about how little they stick out. They're climactic battles that are for the purpose of beating up a big bad and not much else. Careful button timing and smart Battle Stamp (equippable passive buffs) preparation makes quick work of these fights. *Costume Quest* isn't a game made to wear you out or intimidate you with scary monsters; it's much more interested in maintaining a smooth ride throughout so you can sit back and enjoy the scenery.

Playing *Costume Quest* again now has convinced me that more well-known developers should put resources into projects like this. Obsidian Entertainment put

a small team of open world RPG veterans on a tiny RPG game with *Pentiment*, which is analogous to *Costume Quest*. These games don't aim absurdly high, they aim to refine the familiar systems they have within a framework that reflects who made it. It's a little cliché to say, but you can feel the care that went into *Costume Quest*, and you can see where its developers drew the line. Collecting candy, costumes, stamps and buddies, is all it set out to do. There's almost nothing in *Costume Quest* that strikes me as an afterthought.

MECH-OVER

From cardboard to metal



HOMEMADE

- 1 Cardboard helmet
- 2 Polystyrene antenna
- 3 Roller skate shoes



IMAGINATION

- 1 Titanium wings
- 2 Heat-seeking missiles
- 3 Armoured fists

Boo beats

I'm ashamed to admit that I wasn't really a game soundtrack appreciator until relatively recently. *Costume Quest*'s soundtrack is the key to most of its spooky aesthetic and I regret not paying attention to it before.

Double Fine legend Peter McConnell turns what could have been a goofy simplistic soundtrack into something relaxed and layered. The main theme is full of strings and sounds like the stinger for a Halloween episode of a cartoon. Most of the songs go for a 'less is more' approach. The combat music in particular is like a drum-heavy take on Pokémon battle themes. It builds tension by having horns and cymbals simmer in the background, letting the game's colourful sound effects punctuate the open space.

It's such a tiny thing that you might not even hear while trying to smash button prompts, but it supplements *Costume Quest*'s mood, but the mall music might be the best in the game. Late in *Costume Quest* you infiltrate a mall for more candy. McConnell's score sounds exactly like the



Companion characters are just as charming as the main siblings.

Above: Only in *Costume Quest* can the Statue of Liberty and a mech beat up an alien.
Right: You can purchase Battle Stamps that passively increase your stats.



Christmas songs that come pouring out of the crappy speakers in places like that. The piece sounds like the build-up to a hit single for the entire length. It steps out of the way and mirrors what it's like to push through the rush of people in late December while you're trying to get last-minute presents for your loved ones. Even though it's not tense or scary at all, like most of the other songs, it emulates a distinct holiday feeling with surprising precision.

Homemade

Costume Quest runs near flawlessly on the Steam Deck, the ideal place for small games like this. The handheld's small screen helps make up for the game's fuzzy textures.

Double Fine built this game alongside bigger projects and it feels like some of that influenced how the game can fit into a busy life. When I first played *Costume Quest*, the atmosphere absorbed me, and I think I powered through it in one sitting. Now, in between work and playing other games, I can treat it more like a quick burst of fall vibes. It's almost like

hearing Christmas songs on the radio. A quick jaunt over crunchy leaves and discarded candy wrappers and you're immediately in the Halloween mood, then you close the game and hold onto that feeling until you need another hit.

It's a little sad to me that *Costume Quest* didn't receive any sort of high-definition update. I can appreciate how the game looks even on the Steam Deck, but that might not be the same for newcomers. The game's bright, cartoon-ish characters could use a little help. Videogame publishers are remaking fantastic-looking games like *The Last of Us* and *Horizon Zero Dawn* and games like *Costume Quest* are left behind. Given the influx of wholesome and cosy games and gaming's growth in the last few years, I'd like to think there's an audience of witch-lovers and skeleton-fanatics that would fall in love with it.

Costume Quest's modest design surprised me 12 years ago, and now it's simply refreshing. There was a moment while playing *Elden Ring* this year when I passed into Altus Plateau. It's an orange and yellow section of The Lands

Between that screams autumn. I loved it, but I also wished there was more to do there. *Elden Ring* is pretty empty when it comes to other people to speak to or help out. It's a game about a hostile world and surviving all the monsters within it. It wasn't long into exploring Altus that I thought about *Costume Quest* and how valuable it is to have games that dedicate themselves to one distinct aesthetic.

More games should look at how *Costume Quest* trades an obsession with genre, and the expectations that come with that, and puts the most work into its setting. *Costume Quest* is a totally acceptable turn-based RPG, but it's a wonderful Halloween game above all.

It blends art style, music, and quippy writing to illustrate a world filled with jack-o'-lanterns and cobwebs. Your only goals are to roam around in your costume, collect candy, uncover a mystery, and save your sibling. It's not a job for heroes or guns; it's the kind of journey a kid would imagine, and that's why it remains one of my favourite games to play. ■





RETRO

The tiny palmtop with big ideas: Psion Series 5

In his quest for the perfect portable writing device, David Crookes goes back 25 years to put the Psion Series 5 through its paces.

Over the years I've used many word processors, from Protext on the Amstrad CPC and Pages on the Mac to iA Writer on the iPad and Word on a PC. I've inputted many paragraphs into Google Docs and enjoyed writing books using Scrivener. Yet I invariably return to LibreOffice on my Mac, primarily because I like the comfort of a desktop computer and I have a penchant for open-source software.

At the same, however, I've long looked for a way of writing on the go that doesn't involve lugging a large laptop around with me and gives me a proper, physical keyboard. Regular readers may recall that this hunt led me to try the Amstrad NC100 Notepad – a portable Z80-based computer that impressed me. But it was not without its foibles and, today, I'm sad to report that it's been fired and currently sits in a drawer, its tiny screen ultimately proving to be uncomfortable.

Although I could simply pair an iPad with a keyboard (and I have, admittedly, done just that), I've also recently become

acquainted with the British-made Psion Series 5. Making its debut in 1997, it's akin to a miniature laptop – one not much larger than a spectacles case – and my time spent with it has been like the first flush of love. It soon became abundantly clear why this appealing little device is judged to be iconic.

For those with a yearning for history, the Series 5 is the follow-up to the Series 3. There was never a Series 4 – Psion was concerned about tetraphobia in its Asian markets. But in making its next palmtop, it also leapt forward in other ways. This was immediately apparent when prising open the lid and finding it hid a little bit of magic.

Upon lifting the screen, it rotated into position, coming to rest at what appeared to be a perfect angle. At the same time, the beautiful-looking Qwerty keyboard slid forward in a controlled manner, protruding

Above: The built-in apps are user-friendly and very well designed.

"For those with a yearning for history, the Series 5 is the follow-up to the Series 3. There was never a Series 4 – Psion was concerned about tetraphobia in its Asian markets."

from the front of the case,

When it settled, the device felt as if it had expanded in size and it was mightily impressive – a testament to the work of industrial designer Martin Riddiford. The quality of the keyboard also becomes immediately apparent. Looking like one typically found on a

laptop, there are no gaps between the responsive keys, some of which also boast a good number of secondary functions.

Glaring issue

One of the (many) great things about this small-factor machine is that it can run independently of a computer. It's also powered by a couple of AA batteries, so there are no worries about

finding replacements, and they'll keep the machine going for up to 35 hours. That's a staggering figure given how long (or not) modern devices tend to stay active.

There's also a slot for a backup Lithium CR2032 battery, ensuring your data isn't lost when you swap the main batteries, and you can also power the device from the mains, which is reassuring. With the power on, however, the first problem becomes apparent. The 16-colour, 640 x 240, greyscale, 5.6in, monochrome resistive touchscreen is hard to see.

This caught me by surprise. A screen, after all, is rather important and yet, unless you're in a very well-lit room, it's a struggle to catch what's being displayed – and that's before you take into account the effects of glare. Granted, the Series 5 fared much better when used outside in direct sunlight, but winter is coming so that doesn't bode well. It took a fair bit of positioning – both of the Series 5 and my head – in order to see the screen with any great clarity, and I very nearly gave up.

Eventually, though, I discovered that the screen became clearer if I pulled it towards me about ten degrees, and this made me wonder whether Psion had got the angle wrong or whether I was using it wrong by having it flat on the desk rather than in the hand. But I wanted to see if there was another solution. Cue a flick through the 214-page manual and a suggestion that pressing the Fn key and tapping either "M" or "P" would adjust the contrast. The result of this, in truth, didn't have me leaping around the room with joy.

I had better luck with the backlight. Tapping Fn and the spacebar at the same time made the screen glow and it was a vast improvement, especially in dim light. Trouble is, it also made a very annoying buzzing sound and it kept turning off after a short amount of time. To make matters worse, the backlight quickly drained the batteries, too.

There was, I surmised, every chance that this was due to the device's age but, no, that proved not to be true. Reviewers in the late 1990s expressed the same frustration (both of the poor screen and the sound). To its credit, Psion listened, making slight improvements to the follow-up Series 5MX a couple of years later. Its backlight doesn't buzz nor drain the batteries



quite so much, and I kicked myself for not having researched this earlier.

The Psion Series MX also upped the standard memory to 16MB (Series 5 users had a choice between 4MB and 8MB, the latter on test here) and it doubled the processor speed – the Series 5 runs a 32-bit RISC-based ARM 7100 CPU at 18.432MHz, with the MX being a 710T CPU running at 37MHz.

And yet that's not to say the original has felt sluggish. Quite the opposite, in fact. Once it became possible to see what was happening on the screen, I could witness the Series 5 moving at a fair old pace. I was also quickly charmed – here was a device that felt user-friendly and efficient. Psion had put more than 100 engineers to work on this system over two to three years, and the quality still shows to this day.

Key to success

But then Psion was a company to be admired. Formed in Britain in 1980, it had developed software for Sinclair computers such as the ZX81, ZX Spectrum and QL. In 1984, the firm launched its first hardware product, the 8-bit Psion Organiser, and that led to a successor, the Psion Organiser II, a couple of years later.

By the late 1980s, Psion was making its biggest strides, starting with the development of an operating system called EPOC, which was released in 1989 (the name either means "epoch" or "Electronic Piece of Cheese"; no-one has said which for sure). This closed operating system was first used in a series of notebooks. It was also installed on the Psion Series 3 range from 1991.

Series 5 represented a new beginning, though. For this

The Psion Series 5 launched in 1997.

Opening and closing the Series 5 wears out the screen cable.



device, the OS had become EPOC 32, with Psion also deciding it would not be backwards-compatible with the Series 3. It ushered development into the 32-bit era with a PDA that didn't need to make a bunch of compromises. The biggest problem was persuading Series 3 users to upgrade, but there were many reasons to do so.

The excellent word processor became one of them. It can be launched by tapping on the Word icon that's permanently etched at the bottom of the screen (either using a finger or the stylus bundled with the device). At this point, the Series 5 displays a blank file (unless you've already begun work on a document). It also means you can finally put that keyboard to the test.

First impressions were favourable. After bashing down on a few keys, I noticed they had a lot of vertical travel and, while this meant having to press a little harder than normal, causing me to miss the odd key here and there, I soon got up to a decent speed (about half my normal rate). There have been claims that you can touch-type on this keyboard, but I couldn't manage it comfortably. I did, however, become rather adept with the stylus.

Although there are plenty of keyboard shortcuts, the stylus is a beautiful slither of plastic and metal, designed so that you have to hold it away from the tip. It means you can see more of the screen as you poke and prod, and it's proven perfect for stabbing at the various onscreen menus to select bullet lists, headings, different-sizes fonts, bold, italics, underline and paragraph styles.

I even found myself reaching for it to start a spellcheck, keeping it in hand to ignore or correct misspellings, or else add unrecognised words to the dictionary. The stylus really came into its own when inserting objects and it's possible to add sketches, whether drawn freehand or using shapes and other effects. These features extend to the Sketch app, and it was fun to doodle and play around with clip art.

It's also possible to cut, copy or paste, and zoom in and out of



the screen. What's more, there's an excellent spreadsheet program as well. Called Sheet, this app, like the word processor, has been very well considered by the developers and you can use it to input data and create embedded graphs in your documents. In fact, you can see a clear advantage in Psion's approach to the Series 5: the hardware and software feel tightly woven. They really are made for each other.

Setting standards

The more I delved into this device, the more I became smitten. I tried to put myself in the shoes of someone coming across this in 1997 – at a time when Windows 95 was still ruling the roost and when Apple had only just released Mac OS 8 – and I could see just why the reviews at the time were so gushing. Folk were claiming it had set the benchmark for the mobile office, and there's a strong element of truth in that.

There's just something about the Psion Series 5 that has kept me coming back these past few weeks and I've come to love the thinking behind the built-in applications. Data, for example, is a useful app that lets you create contact cards featuring

The Series 5 benefitted from lots of software and accessories.

"Although there are plenty of keyboard shortcuts, the stylus is a beautiful slither of plastic and metal, designed so that you have to hold it away from the tip."

names, addresses, phone numbers, email addresses and more. It allows you to create separate files for work, home and other activities, while offering a find feature so you're not forever moving up and down lists to locate the information you need.

Calc is another great app to have at your fingertips and, if you think this would be an obvious addition, you've not been using an iPad. Psion saw fit to include a desktop and scientific calculator, even allowing you to flip the layout if you'd rather have the buttons on the left (you can use keyboard shortcuts for the numbers and functions, too).

But arguably the best built-in app is

Agenda, for keeping track of meetings, events and so forth. Boasting a skeuomorphic design, it remains one of the best examples of a time manager I've seen thanks to its inherent simplicity married to an unexpected level of depth. Tap a date, hit Sketch and you can handwrite your entry. I also found it easy to quickly create to-do lists and instantly see if an event was going to stretch into the next day simply by looking for the appropriate icon.

Of all the built-in goodies, however, tapping the Extras icon really unleashed my curiosity. Nestled within was an entry called Program that allows users to code using the built-in Open Programming Language (OPL) – one that's part of the BASIC family. I was initially disappointed to see this feature had only been given a couple of pages in the manual, but I then saw mention of a set of instructions on the

PsiWin CD-ROM.

Hello World

PsiWin, I should explain, was the connectivity software used by Psion users to transfer their files to a computer. You can still download a version of it today (tinyurl.com/APC515PSION), if your device doesn't come with the CD-ROM and, if you grab yourself a serial-to-USB adaptor to connect an RS232 serial cable to a

PC, you'll be able to make use of it on older PCs.

It's better to use a FAT16-formatted CompactFlash Card to move files to a modern computer, however, and this, I gather, is what most people now do. But what of that set of coding instructions? Well, once, again the internet comes to the rescue (tinyurl.com/APC515PSI) for those without the CD-ROM and it allows you to get to grips with a programming language that Psion users had taken to rather well.

For what it's worth, I tried a quick PRINT "Hello World", at first, typing the simple instructions in between the keywords PROC and ENDP, translating the module and creating a program in a format Psion could run. I then looked to create something a little more elaborate and, while I didn't finish any project I began, I nevertheless found the excellent language to be rather versatile.

Even so, I was also intrigued to see just how many commercial, shareware and public domain programs had been created for the Psion. While there isn't a flourishing community of Psion Series 5 coders today (at least as far as I could see – please let us know if there are!), one look at an A5 brochure included in the sale of my device made for fascinating reading. Programs were mainly aimed at business users and there's lots to go at.

Presentations, for example, could be created using Presentation Pack (0/10 for the



The screen can be tricky to read, but fiddling with contrast helps.

name). Self-assessments could be made easier using Shark A Tax (10/10 for that name – someone should resurrect it).

You could buy a personal finance manager called Pocket Quicken, snap up an AA travel companion or invest in a game of Monopoly. Or you could just stuck with the built-in game, Bombs – a clone of Minesweeper that whiled away an hour or so one idle wet Tuesday afternoon.

Find these third-party applications online and you can pop them on your CompactFlash Card and still enjoy them today. But it's the core experience that Psion created that continues to capture my imagination.

It's just a shame Psion didn't stay in the consumer electronics business for much longer – it

exited in 2001 – and some of that was a consequence of Microsoft licensing a handheld operating system, Windows CE.

For, in the eyes of many, Psion had simply become the outsider; one producing devices costing more than its rivals. Yet it had already decided to form Symbian Ltd together with Nokia, Ericsson and Motorola in 1998, and that resulted in EPOC evolving into the Symbian operating system that became widely used on smartphones for the following nine years.

While that could be seen as Psion's legacy, the Psion Series 5 continues to inspire: consider Planet Computers' Gemini PDA and Astro Slide. As it is, I think I'll be exploring the Psion Series 5 for some time yet. ■

The stylus is a great way to navigate menus and draw on screen.





HEATBIT CRYPTOCURRENCY MINING HEATER

Money may not grow on trees, but you can now mine it with a heater.

There's no shortage of articles detailing how proof-of-work cryptocurrencies are so computationally intensive that they are labelled as extremely environmentally unfriendly, but if you could use all that wasted thermal energy to heat your house in winter, the equation wouldn't seem quite so bad. Bitcoin's well beyond the threshold of making a mining heater financially worthwhile, but if you're willing to dabble in more obscure Alt-coin mining you may actually be able to find a way to make a profit by heating your house with these \$1,149 devices. Even if you do, it's still not great environmentally since Australia's grids run predominantly on coal.



POPULAR WAR SIM REGULARLY USED TO MISREPRESENT REAL CONFLICT

Player Unknown's Battleground foundational game used to mimic real combat.

The makers of the popular war simulation game *Arma 3* released a blog post detailing how its in-game footage keeps being used for fake news on the war in the Ukraine. *Arma 3* is a military simulation game released by Czech republic based developer Bohemia Interactive in 2013, and while that's now a pretty old title, there is so much desperation for information and footage on the situation in the region that people are pixelating and darkening footage to make it look like plausible war footage. The game has been misused to represent conflicts in Syria, Israil, Afghanistan and Pakistan over its lifetime.




CHIP CHAT

You can play Microsoft Flight Sim in Microsoft Flight Sim

We wonder how long till someone puts Doom on the cockpit computer.

The 40th anniversary update of the most recent iteration of *Microsoft Flight Simulator* includes a throwback easter egg that allows you to play the first four versions of the game from a cockpit computer. In order to fire up these legacy titles you'll have to load yourself into a DA62 plane, enable cockpit view, turn on the Emergency Locator Transmitter and choose between *Microsoft Flight Simulator 4.0*, *3*, *2*, or *1*. These games are fully playable with a cursor and set of on screen buttons to let you fly planes from inside another plane on the tarmac.



OPENAI MAKES AN ALGORITHM THAT'S BETTER THAN YOU AT MINECRAFT

Diamond tools and the key to much faster AI training.

OpenAI has trained an AI algorithm to learn in a new way that will dramatically change the range of applications and effectiveness of machine learning. The new technique is referred to as imitation learning, and was made possible by consuming 70,000 hours of *Minecraft* footage that allowed the algorithm to build the skills required to perform complex tasks like crafting diamond tools. The main breakthrough came from training another AI to label the 70,000 hours of footage after being trained on a 2,000 hour dataset with click logs and labelled outcomes, compiled by researchers. This more accurate AI offers the potential to complete abstract tasks like controlling robotic limbs and driving automated vehicles when ported to other applications.

PC GAMER

THE GLOBAL AUTHORITY ON PC GAMES

**Over 22 million gamers
visit us online every month.**

Join them.

www.pcgamer.com

SAMSUNG

Smart Monitor M8

Work, watch and live in style.


apc
EDITOR'S
CHOICE



Experience Monitor and Smart TV all in one. Access your content, entertainment and streaming apps – all without connecting to a separate PC. Featuring an ultra-slim design and detachable camera, the Smart Monitor M8 delivers a new take on connectivity and productivity.



Smart TV experience¹



Iconic slim design



Magnetic SlimFit camera



PC-less workspace²



IoT Hub³

samsung.com/au/smart-monitor

¹App subscriptions may be required. TV tuner is not included. ²Requires internet connection. ³Requires a Wi-Fi, Bluetooth or other wireless network connection and devices registered to SmartThings app.